Remote Sensing of Clouds with Longwave Infrared Cameras at the Pierre Auger Observatory

Patrick van Bodegom

School of Physical Sciences
University of Adelaide

This dissertation is submitted for the degree of
Doctor of Philosophy

May 2018

Contents

Abstract

Knowledge of atmospheric conditions at the site of a cosmic ray observatory is important, especially for measurements made using the fluorescence technique. At the Pierre Auger Observatory in Argentina, an extensive network of meteorological instruments is dedicated to atmospheric monitoring, several of which are used for the remote sensing of night-time clouds. Clouds can be identified passively by detection of the thermal radiation they emit, and are often strong radiators at long-wave infrared wavelengths. As part of the University of Adelaide's contribution to the Observatory, we have installed four infrared cameras at the Observatory for cloud detection. The current generation cloud cameras are radiometric, and are sensitive to the 8–14 µm waveband. However, identifying clouds is not necessarily straightforward as atmospheric water vapour also absorbs and emits radiation at these wavelengths.

In this dissertation, I present the method that I use to identify clouds in our thermal images. Another major focus of my studies has been to calibrate our cameras. However, as they were already collecting data at the Observatory, the routines had to be developed remotely. These methods have been reproducible for each of our cameras, and could perhaps benefit other researchers in this field.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Signed: Date: 18/05/2018

Acknowledgements

I would like to thank my supervisors, Roger Clay and Bruce Dawson, for the help and wisdom they have offered me over the course of my PhD. Their support and enthusiasm have been motivating, and without their guidance this would not have been possible. Additionally, I would like to thank Jose Bellido and Gary Hill for their assistance and input provided in weekly discussions.

I would also like to thank all members, past and present, of the High Energy Astrophysics Group in Adelaide. It has been wonderful to be surrounded by such a friendly and helpful environment. In particular, I would like to thank Phong, Simon, James, and Alex for their advice and friendship. These past 8 years with them have been righteous.

I would also like to thank Neville Wilde for his technical support, Trent Grubb for maintaining the operation of the cloud cameras and for his help over the years, and Jarryd Day and Padric McGee for their frequent input in cloud camera discussions.

I am grateful for the opportunity to have been involved with the Pierre Auger Observatory. I would like to thank all collaboration members for their input and the resources provided during my research.

My family and friends have provided me with much support, especially over the past year. In particular, I would like to thank my parents, Ben and Kerri, my sister, Claire, and my broham, Kiefer. Without their on-going support (both emotional and financial) and encouragement, I wouldn't have made it this far.

Last but not least, I would like to thank Rhianna. Her love and companionship have kept me sane, and her belief in me has kept me going.

Nomenclature

T_b	Brightness temperature
T_{cam}	Temperature of the camera's sensor
T_{dew}	Dew point temperature
ADC	Analogue to Digital Converter
Auger	The Pierre Auger Observatory
CLF	Central Laser Facility
FD	Fluorescence Detector
GDAS	Global Data Assimilation System
GOES	Geostationary Operational Environmental Satellites
MODTRAN	MODerate resolution atmospheric TRANsmission
TPWV	Total Precipitable Water Vapour
XLF	eXtreme Laser Facility

Chapter 1

Cosmic Rays and the Pierre Auger Observatory

The Pierre Auger Observatory [1] is located in the Mendoza Province of Argentina and has been designed to study the most energetic particles in the Universe (for cosmic rays, see Section 1.1). Built upon the vast plain of Pampa Amarilla at the base of the Andes mountains, the site covers an impressive 3000 km^2, making it the largest observatory of its kind.

When a cosmic ray interacts with the Earth's atmosphere, it initiates a cascade of particles (an extensive air shower) which can be observed with specialised detectors. In this sense, the atmosphere acts as a detection medium. Consequently, the properties of the medium (such as the atmospheric conditions) must be well understood. To this end, an extensive network of meteorological instruments operate at the Observatory (see Chapter 3).

Clouds can adversely impact upon some of the extensive air shower measurements (see Section 1.2.3). As part of the University of Adelaide's contribution to the Observatory, we have installed four infrared cameras distributed over the site for night-time cloud detection. This chapter provides some context for my research. Background information that relates to how our cameras function may be found in Chapter 2 and Section 3.1.

1.1 Cosmic rays and extensive air showers

Cosmic rays are high-energy charged particles of extraterrestrial origin. Since their discovery over a century ago [2], they have been the focus of much research. Although cosmic rays with low to moderate energies are somewhat understood, those at the highest energies (with primary energies beyond 10^{17} eV) remain mysterious. The astrophysical processes capable of accelerating particles to such immense energies, the resulting composition of the beam of particles, and the origins remain major unsolved problems in modern-day physics. What hinders answers to these questions is the extremely low flux at such energies, which for cosmic rays above 10^{20} eV is less than 1 particle per km^2 per century reaching the Earth. This property is described by the cosmic ray energy spectrum.

1.1.1 The cosmic ray energy spectrum

The cosmic ray energy spectrum describes the number of particles arriving at Earth as a function of energy (see Figure 1.1a). It follows an inverse power law, and can be expressed roughly as:

$$\frac{\mathrm{d}N}{\mathrm{d}E} \propto E^{-\gamma}\,\mathrm{m}^{-2}\,\mathrm{s}^{-1}\,\mathrm{sr}^{-1}\,\mathrm{eV}^{-1}\,, \tag{1.1}$$

where E is the particle energy, N is the number of particles arriving at Earth, and γ is the spectral index.

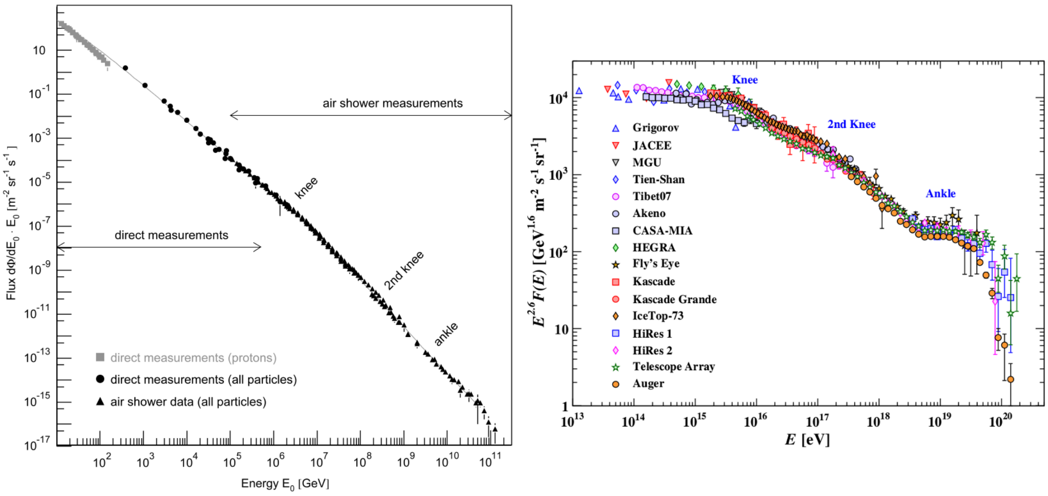

(a) **(b)**

Figure 1.1: **(a)** The measured cosmic ray energy spectrum. The deviation from a single power law can be seen over this large energy range. Anomalous features are indicated [3]. **(b)** The cosmic ray energy spectrum measured by various experiments. The flux is multiplied by $E^{2.6}$ to highlight various features (the knee, the ankle, and the suppression of the cosmic ray flux beyond $10^{19.5}$ eV) [4].

The spectrum covers an enormous range, extending from MeV energies to at least 10^{20} eV. Surprisingly, a spectral index of $\gamma \sim 3$ describes the majority of the spectrum rather well (from about 10^{10} eV to 10^{20} eV). This power law function indicates that the particles are of a non-thermal origin [3].

Upon closer inspection, the spectrum appears to deviate from a single power law (which indicates a variable γ). The anomalous features are more obvious when the cosmic ray flux is scaled by some energy-dependent factor such as in Figure 1.1b. The spectrum first steepens at $\sim 3 \times 10^{15}$ eV (a feature known as the "knee"), and steepens again at about $\sim 10^{17}$ eV (the "second knee") before flattening out at $\sim 3 \times 10^{18}$ eV (the "ankle"). Although there are on-going debates on what causes changes in the spectral index, it is widely believed to relate to the acceleration mechanisms, origins, and propagation of the cosmic rays. In addition, experimental results indicate a suppression of the cosmic ray flux above $\sim 10^{19.5}$ eV (see Figure 1.1b) which is consistent with the Greisen-Zatsepin-Kuz'min (GZK) limit—a theoretical upper limit on the energies of cosmic rays arriving at Earth from distant sources due to interactions with cosmic microwave background photons. This feature was predicted by Greisen [5], and independently by Zatsepin and Kuz'min [6] soon after the discovery of the cosmic microwave background (CMB) [7]. This feature may also be the result of cosmic accelerators reaching their energy limit.

1.1.2 Extensive air showers

Beyond 10^{14} eV the extremely low cosmic ray flux makes direct detection increasingly difficult [8]. Fortunately, cosmic rays in this high-energy regime create particle cascades upon interacting with the Earth's atmosphere (usually with a nitrogen or oxygen nucleus) that are sufficiently large to reach the surface. The cascade of particles is known as an extensive air shower (EAS). As an example, a cosmic ray (the primary particle) at 10^{19} eV can initiate an air shower which develops with $\sim 10^{10}$ particles spread over $\sim 20\,\mathrm{km}^2$ in a thin disc moving at close to the speed of light [9]. A primary particle's properties such as its initial energy, arrival direction, and composition can be obtained indirectly from the air shower it produces. An EAS can be separated into three components:

- **Electromagnetic component** – consists of electrons, positrons, and γ-rays created through a series of electromagnetic particle interactions.

- **Hadronic component** – consists of protons, neutrons, charged pions and kaons.

- **Muonic component** – consists of muons and neutrinos produced through the decay of charged pions and kaons.

Figure 1.2 illustrates the main components of an extensive air shower. Also accompanying the development of an air shower are atmospheric nitrogen fluorescence and Cherenkov light, which are produced as the charged shower particles pass through the atmosphere (as well as water for the latter). The light can be collected using specialised detectors and used to infer the nature of the primary particle. Two well-established detector types used to record EASs are ground (or surface) arrays and atmospheric fluorescence telescopes.

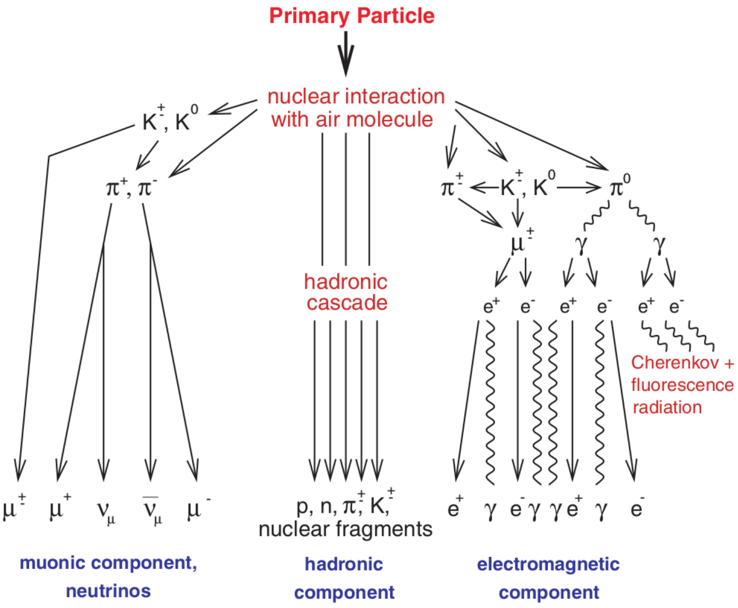

Figure 1.2: A visual representation of an extensive air shower (EAS) initiated by a hadronic primary. The EAS consists of an electromagnetic component, a hadronic component, and a muonic component. From [10].

1.2 The Pierre Auger Observatory

The Pierre Auger Observatory (Auger) is a hybrid detector that employs two independent and well-established methods in order to study cosmic rays through the showers resulting from their interactions with the Earth's atmosphere. Using an enormous collecting area ($3000\,\mathrm{km}^2$), it accumulates a large number of ultra-high energy cosmic ray events with energies up to $\sim 10^{20}$ eV, from which statistically significant conclusions can be made.

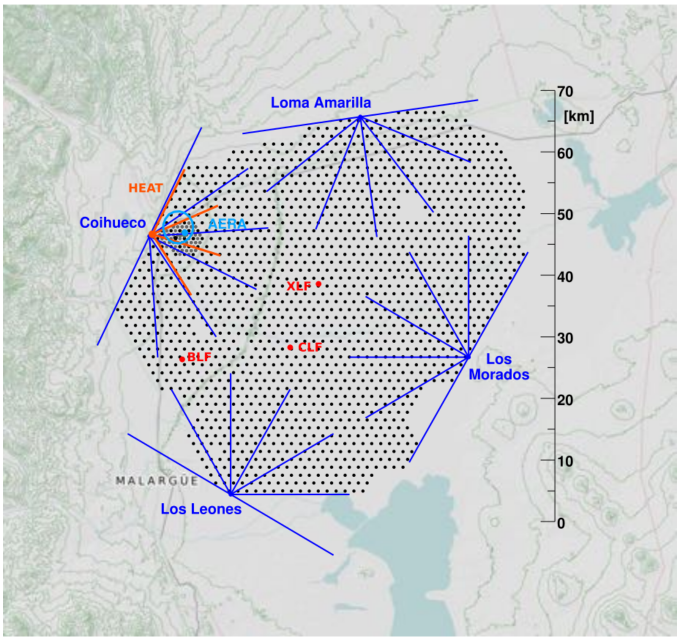

Figure 1.3: The layout of the Pierre Auger Observatory. Black dots correspond to the water-Cherenkov stations. The four main fluorescence detectors are positioned along the perimeter of the Observatory, and overlook the atmosphere above the surface array. Blue lines show the horizontal field of view of each of the 6 fluorescence telescopes at the main sites. From [11].

Figure 1.3 shows the layout of the Observatory. It features a surface detector (SD) array consisting of 1660 water-Cherenkov stations designed to sample EAS particles at the surface-level (about 1400 m a.s.l.). This is complemented by the fluorescence detector (FD) which consists of 27 fluorescence telescopes that detect the nitrogen emission accompanying an EAS.

Since operation began in 2004, data collected with the Pierre Auger Observatory have led to a number of important discoveries in the field of cosmic ray astrophysics [11]. Some of the scientific achievements have included: the measurement of the suppression of the cosmic ray flux above 4×10^{19} eV [12], upper limits on photon [13] and neutrino [14] fluxes at ultra-high energies, the transition in cosmic ray mass composition between 10^{18}–$10^{19.5}$ eV [15], and the anisotropy in the arrival directions of ultra-high energy cosmic rays above 8×10^{18} eV (which supports an extragalactic source for their origins) [9].

1.2.1 Surface detector

The surface detector (SD) is comprised of 1660 water-Cherenkov particle detector stations (chosen for their robustness, large area, and low cost) arranged on a triangular-based grid with a 1.5 km spacing. The detector spacing was chosen as a compromise between energy threshold and cost. With this configuration, the surface array has an aperture of $7350\,\mathrm{km^2\,sr}$ for zenith angles less than 60°, and a 100 % detection efficiency at energies above $3 \times 10^{18}\,\mathrm{eV}$ [16]. An additional sub-array covering a total area of $23.5\,\mathrm{km}^2$ with a smaller station spacing (750 m) is nested within the standard array and forms the low energy extension of the SD.

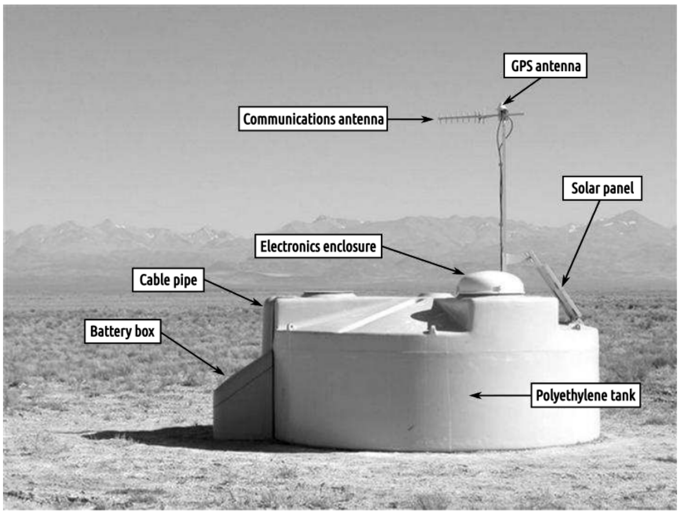

Figure 1.4: A schematic of a water-Cherenkov station at the Pierre Auger Observatory. The main components are indicated. Note that the flatness of the terrain is typical of the site. From [1].

Each SD station is a self-contained unit that operates independently from the rest of the array. Figure 1.4 shows a schematic of one of the stations. Besides the communications antennae and solar panels, the detector components are housed within a polyethylene cylindrical tank with a diameter of 3.6 m. Each tank contains a sealed liner with a reflective inner surface that stores about 12 kl of ultra-pure water (the water reaches a height of 1.2 m). The high purity not only maximises the transmission of Cherenkov light, but it also allows for better long-term detector performance.

As a charged particle moves through a dielectric medium (such as water or the atmosphere), the molecules surrounding the particle become polarised, and emit radiation upon returning to an unpolarised state. When a charged particle moves faster than the local speed of light in the medium $v > \frac{c}{n}$ (with a refractive index n), the emission is coherent and observed as Cherenkov light. Since the refractive index of water is greater than that of air, the energy threshold for Cherenkov production is lower (about 1 MeV in water rather than 21 MeV in air at sea-level for electrons [17]). The stations at Auger use three downward pointing photomultiplier tubes located at the top of the tank liner to collect Cherenkov light created by EAS particles. The recorded signals are digitised by on-board electronics which are described extensively in [1]. The signals, along with timing information (the relative delay between different detectors), can be used to infer the energy and arrival direction of the primary particle.

A major advantage of the surface detector is that it operates constantly, as opposed to the much lower duty cycle of the fluorescence detector (about 15 % [1]). As a result, it can effectively accumulate large datasets. However, unlike the fluorescence detector, it cannot directly observe an air shower's development.

1.2.2 Fluorescence detector

As an EAS propagates through the atmosphere, the charged secondary particles excite atmospheric nitrogen molecules. The subsequent de-excitation of these molecules results in the isotropic emission of fluorescence light between wavelengths of $\sim$ 300–430 nm (see Figure 1.5). Fluorescence detectors (FDs) are specifically designed to capture the ultraviolet nitrogen emission light that accompanies an EAS's development. Hence, they can be used to study the evolution of an air shower in detail.

Fluorescence detectors also provide an almost calorimetric measurement of an air shower's energy, as the energy deposited in the atmosphere by the charged particles is proportional to the number of fluorescence photons emitted [19]. For instance, at room temperature and atmospheric pressure, the absolute fluorescence yield in the main emission band of 337 nm is about 5 photons per MeV of energy deposited by shower particles [20]. Since the emission is isotropic, fluorescence detection is only possible at very high energies [17] when enough light is produced to stand out in the telescopes against background sources. Fluorescence detectors require clear, moonless nights for effective operation. These limitations result in a $\sim$ 15 % duty cycle [1]. The impact that clouds can have on measurements is described in Section 1.2.3.

The fluorescence detector of the Pierre Auger Observatory consists of 24 fluorescence telescopes arranged at 4 sites; Los Leones, Los Morados, Loma Amarilla, and Coihueco, which are located along the perimeter of the array (see map, Figure 1.3). Six telescopes are housed in a specially designed building at each site, and view the atmosphere above the SD. The individual

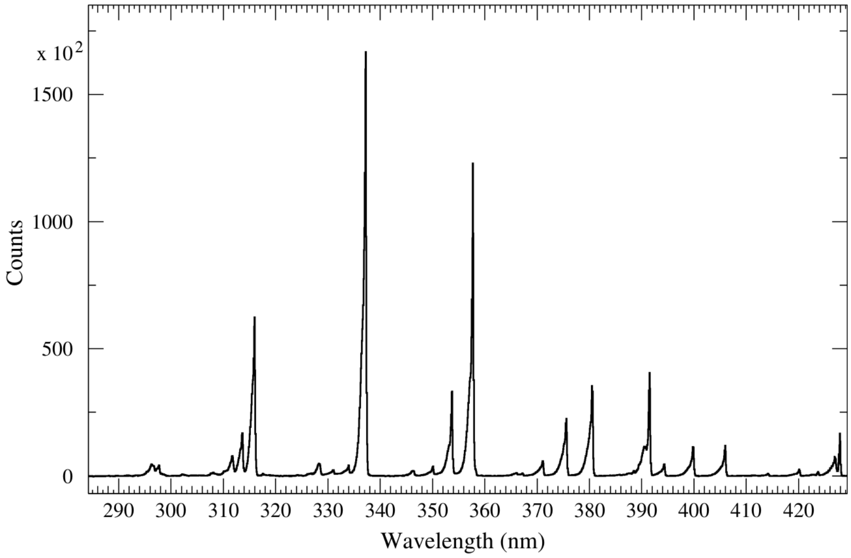

Figure 1.5: The air fluorescence spectrum measured in dry air at 293 K and 800 hPa. From [18].

telescopes have a minimum viewing angle ∼ 1.5° above the horizon, and an approximate 30°×30° field of view in azimuth and elevation, resulting in a combined 180° coverage in azimuth at each site. An additional 3 telescopes comprise HEAT (High Elevation Auger Telescopes), which is the low energy extension of the FD. These particular telescopes are almost identical to those of the standard FD. However, they can be tilted ∼ 30° upwards to observe showers initiated by cosmic rays of lower energies. The altitude at each FD site is about 1400 m a.s.l., with the exception of Coihueco and HEAT at ∼ 1700 m a.s.l.

Figure 1.6 shows a schematic of the Auger FD telescopes. The design is based on Schmidt optics, which was chosen for the reduced coma aberration in large optical systems. Atmospheric nitrogen fluorescence light that passes through the aperture system (and UV filter) is focused onto a specialised camera by a large 13 m^2 segmented spherical mirror. Each fluorescence camera consists of 440 hexagonal photomultiplier tubes (the FD pixels), which are arranged into 22 rows and 20 columns. Surrounding each pixel are 6 light collectors (referred to as "Mercedes stars"), which maximise a camera's efficiency. Details regarding the detector components, electronics, and operation are described extensively in [1].

The FD pixels have a ∼ 1.5° angular field of view. Since clouds can affect measurements made using the fluorescence technique (see the following

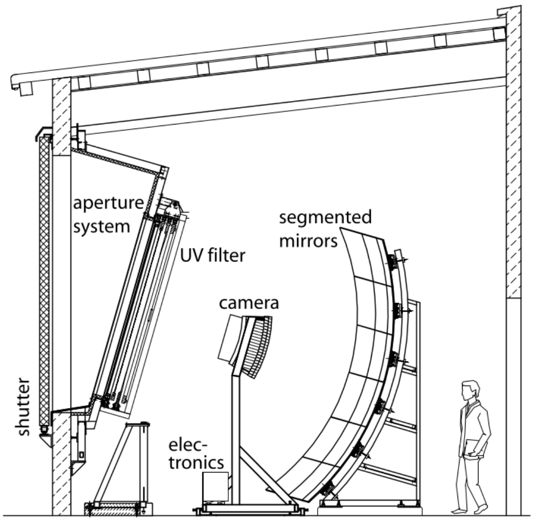

Figure 1.6: A schematic of the fluorescence telescope optical system. From [19].

section), the main objective of our infrared cameras is to identify any clouds within each of the FD pixels' field of view. This information is stored in the cloud camera database, which is one of the atmospheric databases maintained at the Observatory. The procedure is the focus of Chapter 5.

1.2.3 The impact of clouds on fluorescence measurements

One of the challenges with using the atmosphere as a cosmic ray detection medium is that the properties of the system must be known well. Unlike many other areas in physics where experiments can be performed in laboratory environments, atmospheric phenomena cannot be controlled. Although many of the effects that atmospheric conditions can have on extensive air shower (EAS) measurements can be taken into account (for example [21]), it is not straightforward to correct for the impact that clouds can have on observations.

Clouds can affect measurements made using the fluorescence technique. Two scenarios can occur which depend on the location of the cloud [23]. When a cloud layer is positioned between a fluorescence detector and the developing EAS, it can attenuate nitrogen fluorescence through Mie scattering. This is illustrated in Figure 1.7a. As a result, there is a decrease (or

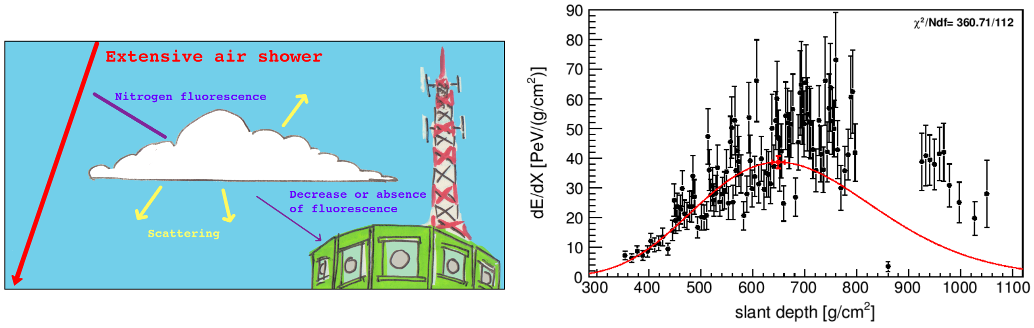

(a) (b)

Figure 1.7: **(a)** A diagram that shows how nitrogen fluorescence from an extensive air shower (EAS) can be attenuated by cloud. Illustrated by Rhianna Skye Carr. **(b)** An example of an EAS's longitudinal profile which has been partially blocked by cloud. Energy deposited into the atmosphere has been determined from the light at aperture. However, nitrogen fluorescence between atmospheric depths of ~ 800–$920\,\mathrm{g\,cm^{-2}}$ has been obstructed by cloud. Black points correspond to the data, and the red curve is the erroneously reconstructed profile. The profile has been taken from [22].

absence) of light reaching the detector which manifests as a dip in an air shower's longitudinal profile (see Figure 1.7b).

When a cloud layer is positioned in the path of an extensive air shower as it develops, the strongly forward-beamed Cherenkov light (usually not directed towards the detectors) accompanying a shower can be scattered towards a fluorescence detector. This is illustrated in Figure 1.8a. As a result, there is an *apparent* excess of fluorescence light reaching a detector which manifests as a peak in an air shower's longitudinal profile (see Figure 1.8b).

Longitudinal shower profiles affected by clouds can lead to erroneous reconstructions of cosmic ray events [23]. The extent of both effects depends on a cloud's scattering characteristics. However, it is not a straightforward task to determine corrections. One of the major challenges is due to the uncertainty in a cloud's scattering properties [24], which depends on its composition, and can be difficult to obtain. Since it is difficult to account for the impact that clouds can have on fluorescence measurements, potentially affected data are often rejected from analyses. For this reason, cloud identification is an important aspect of atmospheric monitoring at the Pierre Auger Observatory. A variety of instruments operate at Auger specifically for night-time cloud detection, many of which are discussed in Chapter 3. Most relevant to my research have been our infrared cameras, which pas-

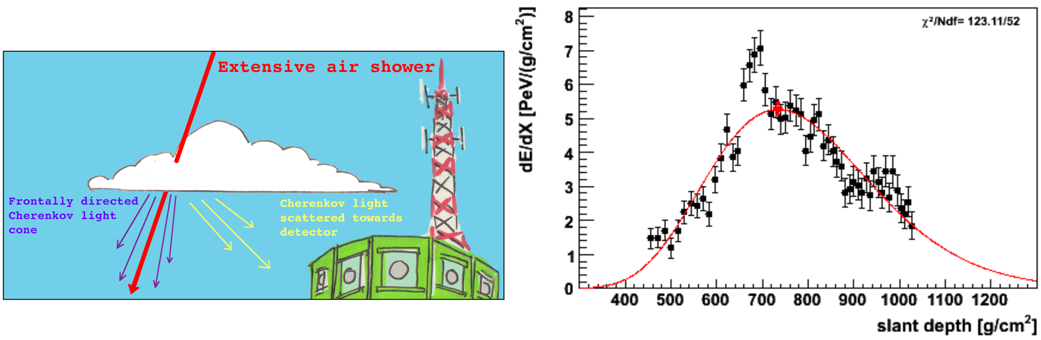

(a) **(b)**

Figure 1.8: **(a)** A diagram that shows how Cherenkov light from an extensive air shower (EAS) can be scattered by cloud towards a fluorescence detector. Illustrated by Rhianna Skye Carr. **(b)** An example of an EAS's longitudinal profile which has been affected by cloud in the path of the shower's development. Energy deposited into the atmosphere has been determined from the light at aperture. However, Cherenkov light at atmospheric depths of $\sim 670\,\mathrm{g\,cm^{-2}}$ has been scattered towards the detector, resulting in an excess. Black points correspond to the data, and the red curve is the erroneously reconstructed profile. The profile has been taken from [22].

sively detect clouds by detection of the thermal radiation they emit.

Chapter 2

Thermal Radiation in the Earth's Atmosphere

The Earth's climate is largely determined by the balance between incoming solar radiation and outgoing thermal radiation. Clouds are known to have a major impact on this radiative balance [25]. Not only can clouds reflect incoming short-wave radiation back into space, they also absorb the Earth's outgoing long-wave infrared radiation. Clouds also emit thermal radiation, and return a fraction of the absorbed energy back towards the ground [26]. This behaviour is also characteristic of certain greenhouse gases, such as water vapour and carbon dioxide, although to a lesser extent.

At night most of the radiation reaching the ground is thermal radiation emitted in the atmosphere [27]. Clouds are particularly efficient emitters, and appear bright in the sky at 8–14 µm wavelengths. This is the basis for many ground-based instruments that can be used to detect clouds, for example a thermal imaging camera.

2.1 Black body radiation

A fundamental concept in physics is that objects at temperatures above absolute zero emit thermal radiation. This is true for clouds, as well as atmospheric gases. Low, thick clouds are particularly efficient at emitting thermal radiation, as their radiative properties at infrared wavelengths resemble those of a black body [28]. A black body is an idealised object that absorbs all incident radiation. Such an absorber is also, by Kirchhoff's law of thermal radiation, a perfect emitter (see Section 2.2.2). The thermal radiation emitted by such an object is known as black body radiation. The spectral radiance of a black body B_λ at a temperature T is given by Planck's law:

$$B_\lambda(T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda kT}} - 1}\,, \tag{2.1}$$

where λ is the wavelength of the emitted radiation, h the Planck constant, k the Boltzmann constant, and c the speed of light in a vacuum.

Black body radiation is emitted over a continuous range of wavelengths, and with a spectral shape that depends only on the temperature of the body. The wavelength λ_{max} at which the intensity peaks is given by Wien's displacement law:

$$\lambda_{max} T = 2900\,\mu\text{m}\,\text{K}\,. \tag{2.2}$$

For a black body at 288 K, a typical temperature for the Earth's environment, the intensity peaks at $\sim 10\,\mu\text{m}$. As discussed later (see Section 2.2), light at this wavelength transmits freely through the atmosphere. For this reason, most infrared sensors are designed to be sensitive to this radiation.

Although a black body is a theoretical construct, the effectiveness of a real body at emitting thermal radiation can be characterised by an additional quantity, its emissivity ϵ. The spectral radiance (or specific intensity) I_λ of a real emitting body can be expressed as:

$$I_\lambda = \epsilon B_\lambda\,, \tag{2.3}$$

where ϵ is some number between 0 (totally reflective) and 1 (a black body). The value depends on the properties of the material and, in general, has some dependence on wavelength. For instance, clouds absorb most radiation (an emissivity near 1) at infrared wavelengths. In the visible spectrum, however, clouds often appear white as a result of scattering sunlight (hence a lower emissivity). These interactions are largely governed by the size distributions of the cloud particles and are discussed in Section 2.4.

It is sometimes useful to express the spectral radiance as a brightness temperature T_b. The brightness at a given wavelength can be characterised by the temperature of an equivalently bright black body. That is, for any value I_λ then T_b is defined by the relation:

$$I_\lambda = B_\lambda(T_b)\,. \tag{2.4}$$

It follows that the brightness temperature is given by:

$$T_b = \frac{hc}{k\lambda} \ln^{-1}\left(\frac{2hc^2}{I_\lambda \lambda^5} + 1\right)\,. \tag{2.5}$$

This can be a useful quantity, as brightness temperatures vary with both the temperature and the emissivity of the emitting surface. Data from weather

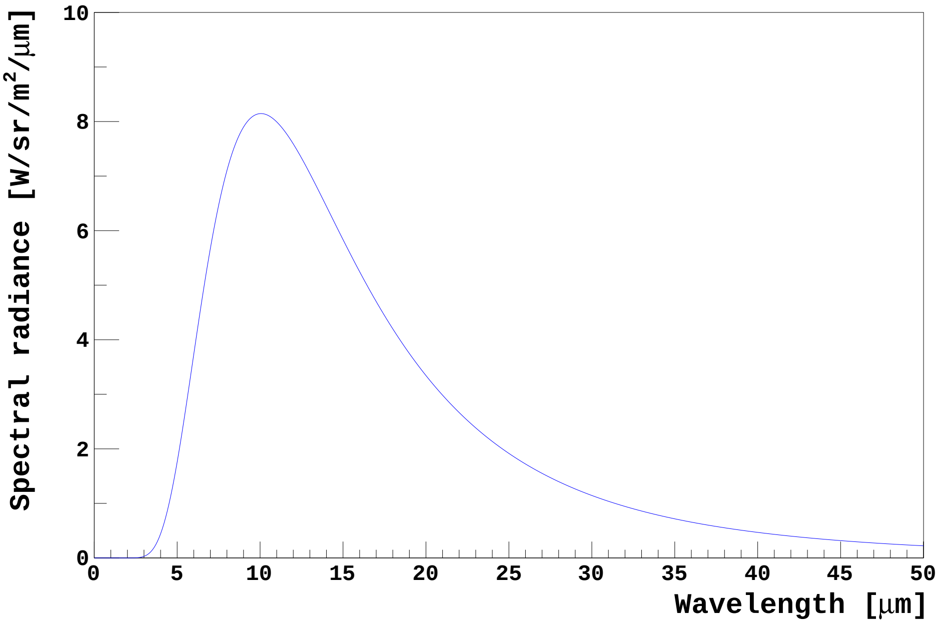

Figure 2.1: The Planck function (Equation 2.1) for a 288 K black body, a typical temperature for the Earth's environment. The spectrum depends only on the temperature of the body. The intensity peaks at $\sim 10\,\mu\mathrm{m}$ (Wien's displacement law, Equation 2.2) and the radiant flux is $\sim 390\,\mathrm{W\,m^{-2}}$ (Stefan-Boltzmann law, Equation 2.6).

satellites are commonly expressed in this way, as upwelling radiance can be used to retrieve surface and cloud top temperatures. An example may be found in Section 3.3.

Another important property of thermal radiation is the radiant flux F, the power emitted per unit area, which rises steeply with the temperature of the emitting body ($\propto T^4$). This is given by the Stefan-Boltzmann law:

$$F = \epsilon \sigma T^4 \,, \tag{2.6}$$

where σ is the Stefan-Boltzmann constant. In principle, an infrared camera measures the power radiated by an object within a given area (a pixel). This is used to infer the body's temperature. However, in practice, a camera is only sensitive to a portion of the black body spectrum (8–14 µm for our model of camera, and with a varying spectral response). Moreover, various camera settings affect the performance of the camera (see Section 3.1.2). For these reasons, to determine a relationship between signal and temperature, it is standard practice to calibrate infrared sensors. This is typically achieved with an experimental black body (manufactured to have $\epsilon \sim 1$) of a known temperature, as its radiative properties are well-defined. Some characteristics of black body radiation are summarised in Figure 2.1.

2.2 Atmospheric transmission

As light passes through a medium, energy may be lost to absorption, gained by emission, and redistributed by scattering. In the context of our research with the infrared cloud cameras, we can consider the medium to be the atmosphere.

The transmission of light through the atmosphere depends on the wavelength of the radiation and, largely, on the concentrations of atmospheric gases (see Table 2.1 for the composition of dry air [1]). The transmittance of light from sea-level out to space for a clear atmosphere has been simulated in Figure 2.2. The extinction of light at wavelengths below $\sim 2\,\mu\text{m}$ is dominated by molecular and aerosol scattering, as well as ozone (O_3) absorption. Determining the optical properties of the atmosphere within the 300–400 nm spectral range of the fluorescence detectors is a major focus of atmospheric monitoring at Auger. However, at the infrared wavelengths observed by the cloud cameras, radiation is most affected by molecular absorption.

Gas	**Symbol**	**Volume [%]**	**Concentration [ppm]**
Nitrogen	N_2	78.08	–
Oxygen	O_2	20.95	–
Argon	Ar	0.93	–
Carbon dioxide	CO_2	0.0403	403
Neon	Ne	0.0018	18
Helium	He	0.0005	5
Methane	CH_4	0.00018	1.8

Table 2.1: The composition of dry air. Not included in the dry atmosphere is water vapour (H_2O), whose concentration varies strongly with both location and time. Table modified from [29] to show the current (October 2017) concentration of CO_2.

---

[1] For the concentration of CO_2, October 2017: `https://www.esrl.noaa.gov/gmd/ccgg/trends/`.

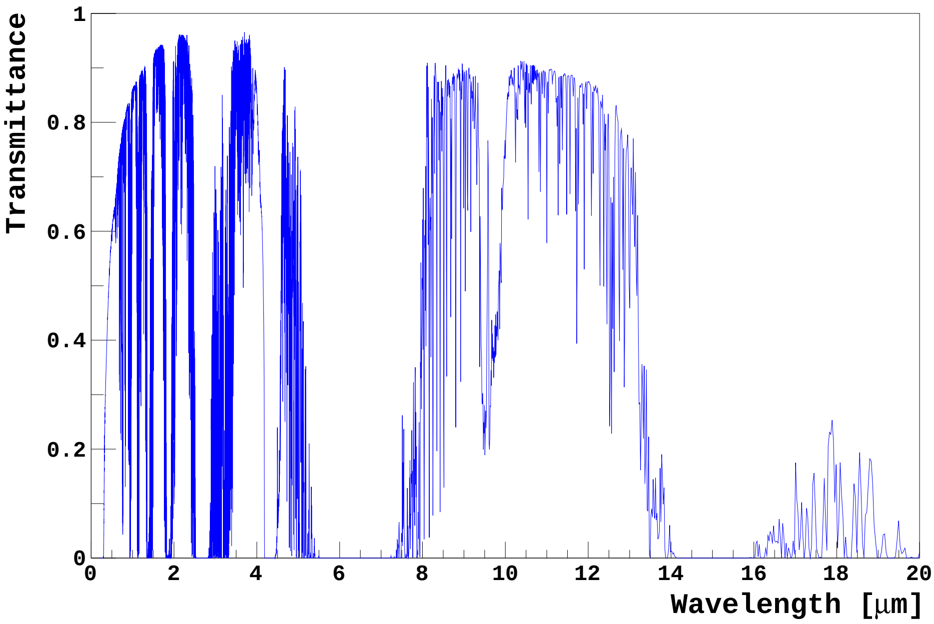

Figure 2.2: Atmospheric transmittance for a clear sky. Simulated by the author with MODTRAN for a 1976 U.S. Standard Atmosphere along a zenith path from sea-level to space. At infrared wavelengths the spectral structure is largely due to molecular absorption by atmospheric gases. Well-known absorption bands are centred at 6.3 µm (H_2O), 9.6 µm (O_3), and 15 µm (CO_2). The region of high transmittance from 8–14 µm is the infrared atmospheric window (IR window).

2.2.1 Molecular absorption in the atmosphere

Molecules can absorb radiation in several ways. One mechanism is for the incident energy to excite orbital electrons of the individual atoms into higher energy states. However, these transitions typically require large energies and are mostly seen at ultraviolet and visible wavelengths. In addition to electronic excitations, molecules can also exhibit rotational and vibrational excitations. Much like the electronic transitions, the energy levels are quantised. The separation between energy levels, however, is often small enough to be excited by thermal radiation.

Not all molecules can directly absorb infrared radiation. There is a strict requirement for there to be a change in the electric dipole moment of a molecule as it vibrates or rotates [29]. This is required for the molecule to couple to the electromagnetic field of the radiation, so that energy exchange may take place [30]. Provided that the frequency of radiation matches the resonant frequencies of the molecule, the radiation is absorbed. For this reason, atmospheric gases such as N_2 and O_2, that do not exhibit a changing dipole moment, cannot directly absorb infrared radiation. Other molecules, for example the greenhouse gases H_2O and CO_2, can.

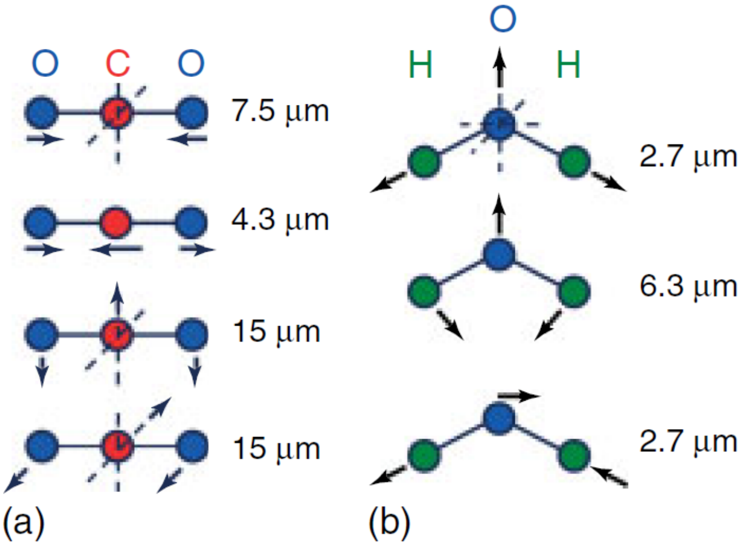

Figure 2.3: Examples of vibrational modes for CO_2 and H_2O. For CO_2, the mode at 7.5 µm has no changing dipole moment. As a result it cannot be excited by the absorption of infrared radiation. However, absorption at 4.3 µm, 15 µm (due to CO_2), and 2.7 µm, 6.3 µm (due to H_2O) can be seen to lower the atmospheric transmittance (in Figure 2.2). Image taken from [29].

Both water vapour and carbon dioxide are triatomic molecules. In general, a molecule comprised of N atoms has $3N$ ways of storing energy (degrees of freedom). This is through translational, rotational, and vibrational motion. For a linear molecule (such as CO_2) this leaves $3N - 5$ vibrational degrees of freedom (there are 3 translational and 2 rotational degrees of freedom). On the other hand, non-linear molecules (such as H_2O) can have only $3N - 6$ vibrational modes (due to the additional rotational degree of freedom). The four fundamental vibrational modes for CO_2, and three for H_2O are depicted in Figure 2.3. In reality, the vibrational modes are often coupled to rotational modes to form rotational-vibrational bands. This leads to the more complex spectra, as seen in Figure 2.2.

For our research with the infrared cameras, the important spectral features are the strong absorption bands centred at 6.3 µm (H_2O), 9.6 µm (O_3), and 15 µm (CO_2). In between the water vapour and carbon dioxide bands is a region of high transmittance from 8–14 µm. This region is known as the infrared atmospheric window (IR window), and allows most of the Earth's thermal radiation to be pass out to space. Most infrared sensors are designed to cover this spectral range, as the passage of radiation is least affected by the atmosphere. Clouds lower the atmospheric transmittance within the IR window (see Section 2.4).

Even for a clear sky there are small features in the transmittance within the window region. These are mostly identified with absorption by atmospheric water vapour [27], although there is also moderate absorption due to carbon dioxide [30]. The absorption between 7–9 µm is due to rotational lines associated with the water vapour rotational-vibrational band at 6.3 µm, while those from 10–14 µm are by groups of pure rotational lines [31]. As the Auger cloud cameras are sensitive to atmospheric water vapour, the amount of moisture in the air must be quantified. Several common approaches are discussed in Section 2.3.

2.2.2 Thermal emission by molecular gases in the atmosphere

For a body to be in thermal equilibrium with its surroundings, conservation of energy requires the amount of radiation absorbed to equal the amount of radiation emitted. This relationship is given by Kirchhoff's law of thermal radiation, usually in the form:

$$\alpha_\lambda = \epsilon_\lambda \, , \tag{2.7}$$

where the spectral absorptivity α_λ and spectral emissivity ϵ_λ denotes the fraction of radiation absorbed or emitted, for a given wavelength λ.

The atmosphere can be considered to be in local thermal equilibrium. To maintain this balance, atmospheric gases emit radiation most effectively at wavelengths at which there is absorption. The intensity of the thermal radiation, however, depends on the temperature of the emitter (as discussed in Section 2.1). This can be described by another form of Kirchhoff's law:

$$j_\lambda = \kappa_\lambda B_\lambda(T) \, , \tag{2.8}$$

where j_λ is the emission coefficient (dimensions of power per unit volume per unit solid angle per unit wavelength) and κ_λ the absorption coefficient (dimensions of inverse length).

Consequently, the strong absorption bands seen in Figure 2.2 lead to several major emission bands. This can be seen in Figure 2.4, showing atmospheric radiance that has been simulated for a clear sky. The most prominent emission features correspond to the absorption bands of water vapour, ozone, and carbon dioxide.

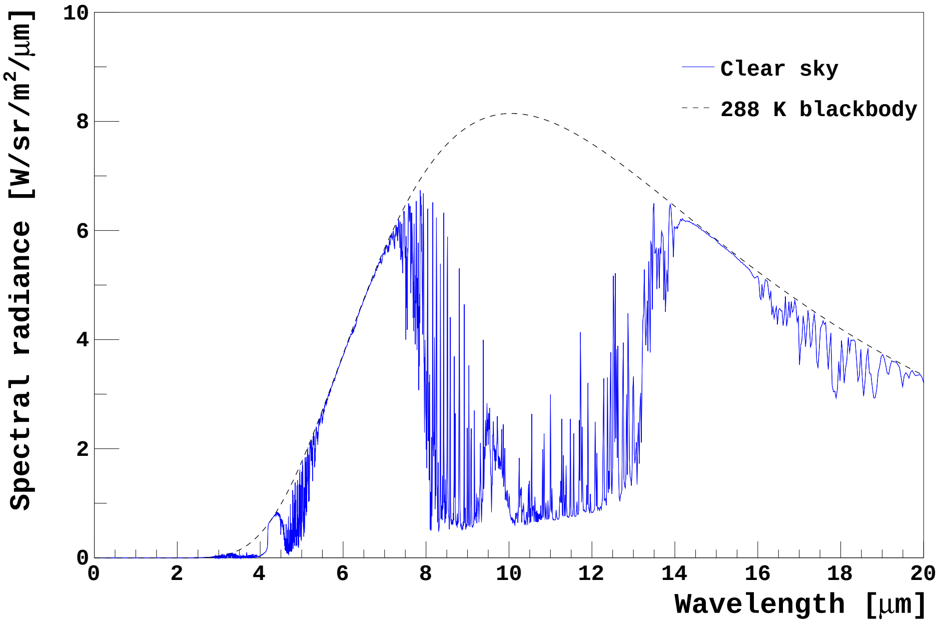

Figure 2.4: Atmospheric radiance from a clear sky reaching the ground. Simulated under the same conditions as Figure 2.2. The well-known emission bands centred at 6.3 µm (H_2O), 9.6 µm (O_3), and 15 µm (CO_2), correspond to the absorption bands (Kirchhoff's law). The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

Owing to the increased path lengths of the main emitting gases (water vapour and carbon dioxide), within the IR window, the sky appears brightest near the horizon [32]. Following the same reasoning, the sky is least bright at the zenith. As a result, it becomes increasingly difficult to distinguish cloud from a clear sky background towards to horizon. To identify cloud within the thermal images from our cameras, it is important to understand the variations in infrared brightness (or effective temperature) with viewing (zenith) angle for a clear sky. The details may be found in Section 5.1.

2.2.3 Radiative transfer

The transmission of radiation through the atmosphere directly affects measurements made by the infrared cameras. Light passing through the atmosphere can be attenuated by absorption or scattering. In addition, the atmosphere also emits thermal radiation. These processes can be formally described by an equation of radiative transfer.

This can be complex to model. To emphasise the difficulty, interactions between radiation and the atmospheric gases are spectrally dependent, with each gas having a unique extinction coefficient. Moreover, the attenuation depends on the vertical distributions of the atmospheric constituents. The presence of aerosols and clouds further complicates this process.

Although simplistic radiative transfer models can be derived and applied [33], more difficult and realistic models must be solved computationally. Fortunately, there are several programs capable of computing transmission spectra within the atmosphere. A particularly well-known one is MODTRAN (MODerate resolution atmospheric TRANsmission), which I have used in this dissertation. Details on MODTRAN can be found in Appendix A.

2.3 Atmospheric water vapour

The concentrations of most atmospheric gases remain fairly constant, or slowly change over time. However, the amount of water vapour present in the air can vary significantly. For the reasons discussed in Section 2.2, our infrared cloud cameras are sensitive to atmospheric water vapour. As an example, this can be seen by comparing an image taken at low humidity (Figure 2.5a) to one taken at high humidity (Figure 2.5b). The decreased visibility in Figure 2.5b results from an increase in water vapour emission, within the spectral range of the camera. This will be further demonstrated in Figure 2.9, showing the simulated radiance for varying amounts of water vapour content. To compensate for this effect, we concern ourselves with quantifying the moisture in the atmosphere.

The amount of moisture in the atmosphere can be expressed in a number of ways. One quantity is water vapour pressure e, the partial pressure of water vapour in air. Air can accommodate only so much water vapour, and the amount depends strongly on temperature T. The air is said to be saturated when the vapour pressure e is equal to the saturated vapour pressure e_s. This can be calculated from the empirical Magnus formula:

$$e_s(T) = \begin{cases} 6.1070 \cdot \exp\left(\dfrac{17.15 \cdot T}{234.9 + T}\right), & T \geq 0\,^\circ\mathrm{C} \\ 6.1064 \cdot \exp\left(\dfrac{21.88 \cdot T}{265.5 + T}\right), & T < 0\,^\circ\mathrm{C} \end{cases} \tag{2.9}$$

where $e_s(T)$ is the saturated vapour pressure (in hPa) for a given air temper-

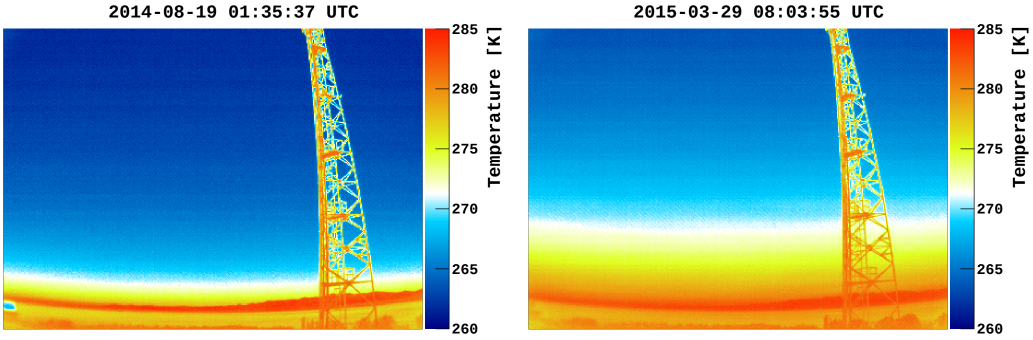

(a) (b)

Figure 2.5: Clear night sky images captured with the infrared camera at the Los Leones fluorescence detector site. At both times, the local weather station measured the air temperature to be 285 K. The relative humidity, however, was recorded to be **(a)** 15 % and **(b)** 70 %. The decreased visibility in (b) is due to increased water vapour absorption and emission.

ature T (in °C). There are various parametrisations of the Magnus formula [34]. To be consistent with the conversion used for preprocessing the GDAS data (for GDAS, see Section 3.6), I use Equation 2.9, which is the same formula as in [21].

The temperature of the air at which saturation occurs is known as the dew point T_{dew}, and defined by the equation:

$$e = e_s(T_{dew})\,. \tag{2.10}$$

At temperatures cooler than the dew point, water vapour will condense. This is an important process for cloud formation [35], and is a criterion I use for estimating cloud coverage in Chapter 6.

Commonly, humidity is quoted in relative terms. The relative humidity u is the amount of water vapour in the air, with respect to the maximum amount of moisture that air can hold. That is:

$$u = \frac{e}{e_s} \cdot 100\,\%\,. \tag{2.11}$$

This is in contrast to absolute humidity, which is a measure of the density of water vapour in air. It can be shown from the ideal gas law that absolute humidity ρ_w is given by:

$$\rho_w = \frac{M_w \cdot e}{R \cdot T}\,, \tag{2.12}$$

where M_w is the molar mass of water, and R the universal gas constant. The relationships between several of these quantities are illustrated in Figure 2.6.

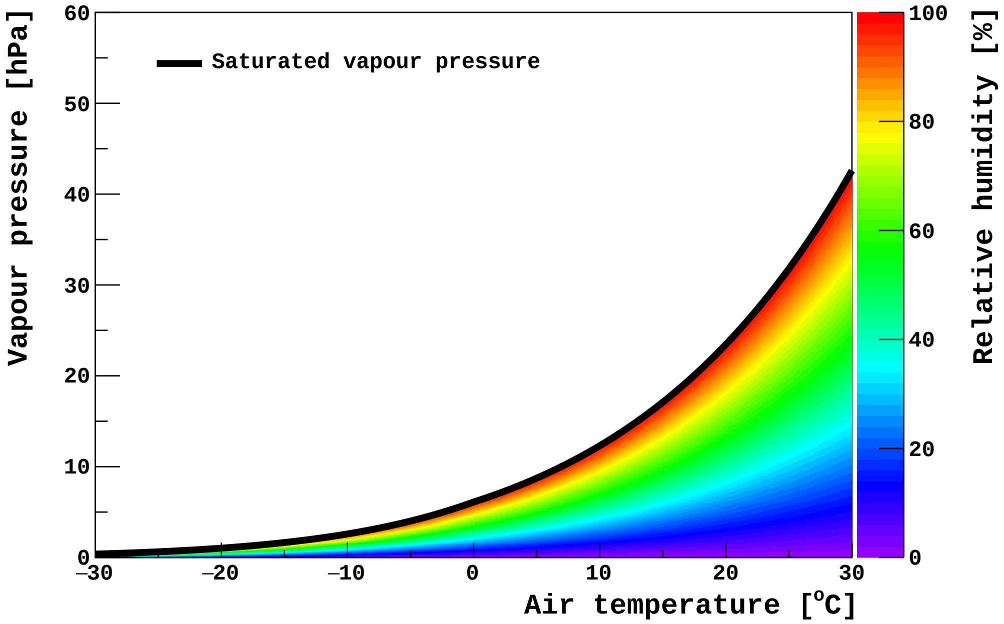

Figure 2.6: The saturated vapour pressure e_s as a function of air temperature T (Magnus formula, Equation 2.9). The relative humidity u gives the percentage of water vapour in the air (Equation 2.11). The air is saturated with water vapour at the dew point temperature T_{dew} (Equation 2.10). Illustrated by the author.

2.3.1 Total precipitable water vapour

The atmosphere can be divided into several layers, which are determined by their temperature profiles. The troposphere is the lowest layer, and characterised by a decrease in atmospheric temperature with height. Almost all the water in the atmosphere can be found in the troposphere [36]. Thermal radiation emitted by water vapour, at all heights, contributes to the infrared flux received at ground. For this reason, it is useful to consider the total amount of water vapour contained within an atmospheric column. This is commonly expressed as total precipitable water vapour (TPWV), the integrated amount of water vapour in the vertical column. Typically, this is determined by integrating a height-dependent profile of absolute humidity $\rho_w(h)$ (Equation 2.12) from surface level h_0 along a vertical path. This is given by:

$$TPWV = \int_{h_0}^{\infty} \rho_w(h) dh \,, \tag{2.13}$$

and the column density (in kg m^{-2}) is often given as the equivalent depth of liquid water (in mm), if all the water vapour in the column precipitated as rain. This conversion is made by dividing Equation 2.13 by the density of liquid water. As the majority of water vapour remains close to the surface (a median scale height of $\sim$ 1.7 km for a similar geographical location [37]), the upper limit in Equation 2.13 only needs to be sufficiently high. I choose to integrate to a height of 20 km a.s.l.

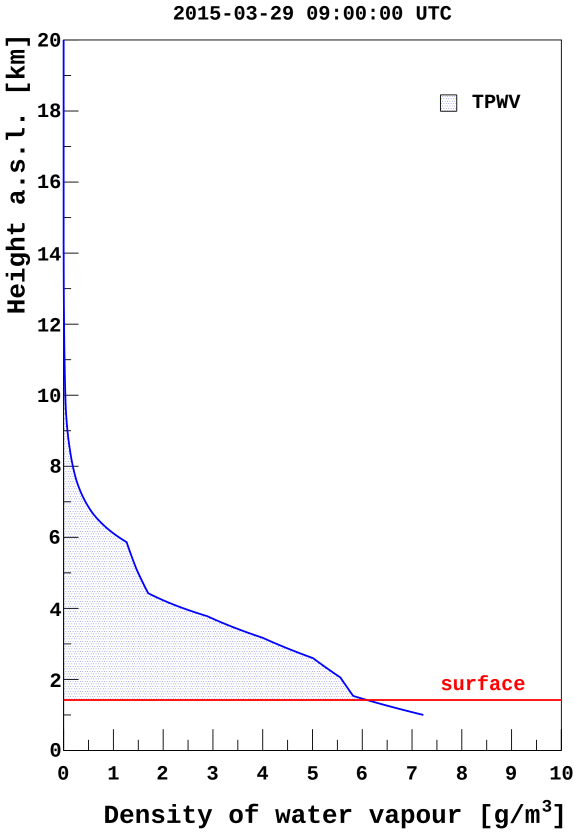

Figure 2.7: An example of using Equation 2.13 to calculate the total precipitable water vapour (TPWV) from an absolute humidity profile. The profile has been derived from temperature and water vapour pressure data (Equation 2.12) predicted by the GDAS model. The atmospheric column (shaded blue) contains $\sim$ 16.1 kg m^{-2} of water vapour, and the equivalent depth of liquid water within the column would be 16.1 mm.

The standard approach to calculating the total precipitable water vapour is to use measured temperature and humidity profiles from a radiosonde. In

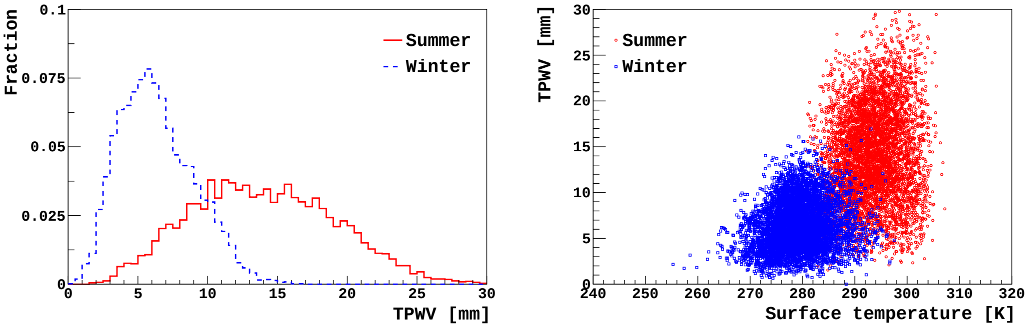

(a) (b)

Figure 2.8: **(a)** Normalised distributions of total precipitable water vapour (TPWV) during Winter and Summer months at the Pierre Auger Observatory. **(b)** TPWV for different GDAS surface temperatures. Both Figures demonstrate that warmer atmospheres can accommodate more water vapour. Determined using GDAS atmospheric profiles from 2009–2016.

the absence of radiosonde data, it is possible to estimate the water vapour content from ground level observations [38][39]. However, such empirical relationships are limited without knowledge of the entire atmospheric column. A common example would be a temperature inversion not being accounted for. Not only would this change the temperature profile, but it can also act to trap water vapour below the ceiling of the inversion. This can weaken correlations made with surface level data [39]. Fortunately, we readily have access to atmospheric temperature and humidity profiles from the GDAS model. As discussed in Section 3.6, the data are available every 3 hours, and globally on a 1° × 1° latitude-longitude grid. This information is implemented into an Auger Offline atmospheric database. An example of using GDAS data to determine the total precipitable water vapour is demonstrated in Figure 2.7.

This procedure can be extended to determine a range of total precipitable water vapours predicted for the Pierre Auger Observatory. Figure 2.8a shows distributions during Summer (Dec–Feb) and Winter (Jun–Aug) months. Generally, the values lie in the range of 0–30 mm. This is typical for mid-latitude locations [40]. It can also been seen there is good separation between the two seasons, with the average water content being lower in Winter (a mean of 6.4 mm and a standard deviation of 2.7 mm) than Summer (a mean of 14.0 mm and a standard deviation of 5.2 mm). Both the increased mean and larger spread during Summer seems reasonable, as warmer atmospheres have the capacity to hold more water vapour. In addition, more evaporation is expected to occur at warmer temperatures. Similarly, this is demonstrated in Figure 2.8b, showing the relationship between precipitable water vapour and GDAS surface temperatures. Although there is some correlation between the two quantities, surface temperature alone is a poor proxy for estimating the atmospheric water vapour content.

One way of demonstrating the effects that water vapour can have on our infrared cameras is to simulate atmospheric radiance for a clear sky, with varying amounts of water vapour. The results can be seen in Figure 2.9 and show nominal amounts of precipitable water vapour in Winter (5 mm) and Summer (15 mm). There is a clear difference within the 8–14 µm atmospheric window, showing that increased thermal emission adds to the overall sky brightness. This is within the spectral range of our cameras, and we are familiar with this in practice (for example Figure 2.5).

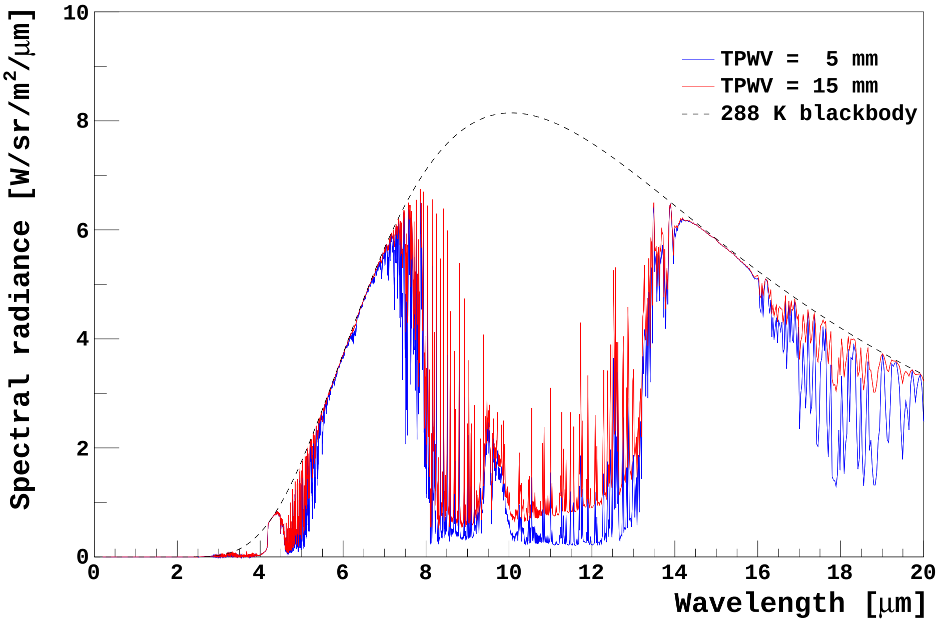

Figure 2.9: Atmospheric radiance from a clear sky reaching the ground. Simulated along a vertical track by the author with MODTRAN for a 1976 U.S. Standard Atmosphere with differing amounts of water vapour. The added water vapour content increases the infrared brightness within the spectral range of our cameras. The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

2.4 Radiative properties of clouds

Atmospheric water vapour can also lead to cloud formation. The majority of clouds are found in the troposphere and can loosely be classified by height as low (surface–2 km), middle (2–7 km), or high (5–13 km) [2]. At low altitudes, clouds are warmer and mostly consist of water droplets. Higher in the atmosphere their composition shifts towards ice crystals.

Clouds display a variety of characteristic shapes and structures, reflecting variations in their development [36]. For clouds to form, moist air must become saturated (a relative humidity of 100 %). This is typically achieved when air rises, and subsequently cools, to its dew point temperature (Section 2.3). For cloud droplets to form and develop into water clouds, there must be a surface onto which water vapour can condense. Atmospheric aerosols can provide such a surface for nucleation to occur [35]. The aerosols that are well-suited for this are known as cloud condensation nuclei. Once formed, a droplet can continue to grow by condensation as more vapour diffuses towards it. This is counteracted by evaporation, instead reducing the size. In addition, droplets can collide and combine to form larger structures. These processes lead to variable growth rates and results in a drop size distribution [36]. This largely determines a cloud's radiative and optical properties [41][42].

The radiative properties of low, thick clouds (such as stratus) resemble those of a black body [27][28]. As a result, they appear opaque at infrared wavelenths and can often be detected by our thermal imaging cameras. Figure 2.10 shows atmospheric radiance reaching the ground, simulated along a vertical track for a clear sky, and a sky with stratus cloud. It can be seen that the emission spectrum of the sky with stratus cloud closely approximates that of a black body (with an emissivity near 1). The intensity of the radiation corresponds to its temperature (Planck's law, Equation 2.1), which is determined by its height in the atmosphere (0.33 km above ground level). The distinction between the clear sky and cloudy sky emission spectra are greatest within the 8–14 µm infrared window. As discussed in Section 2.2, atmospheric absorption at these wavelengths are low. This separation, however, decreases towards the horizon as the path lengths of greenhouse gases are increased.

Although water clouds are typically optically thick at infrared wavelengths, ice clouds (such as cirrus) exhibit rather poor radiative properties [43]. Due to the comparatively low particle density within an ice cloud, cirrus

---

[2] World Meteorological Organisation's (WMO) International Cloud Atlas. The 2017 edition is available online: `https://cloudatlas.wmo.int/clouds-definitions.html`.

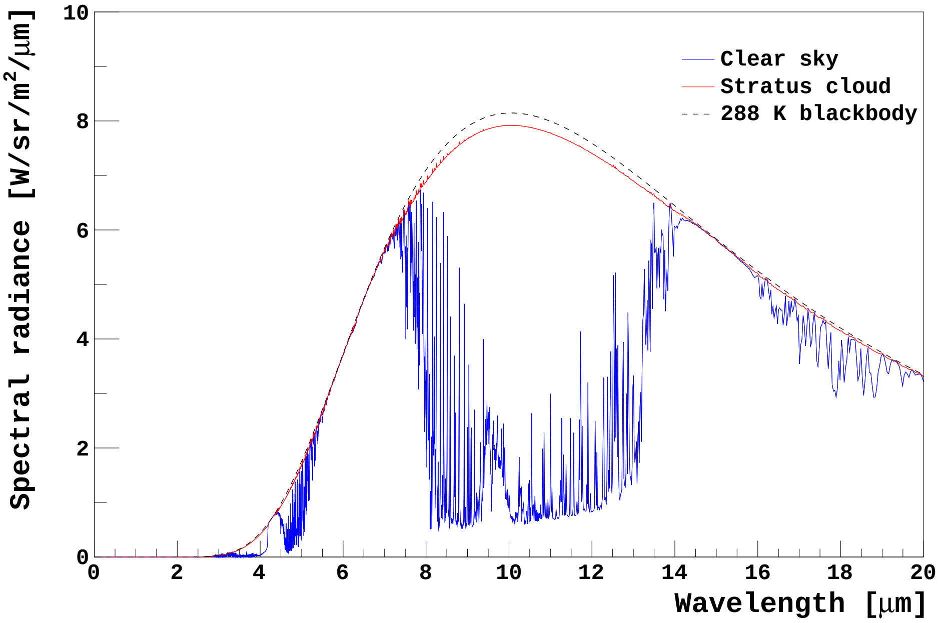

Figure 2.10: Atmospheric radiance reaching the ground from a clear sky, and a sky with stratus cloud at 0.33 km above surface level. Simulated along a vertical track by the author with MODTRAN for a 1976 U.S. Standard Atmosphere. The stratus cloud in this model has an optical depth of ~ 38 at 0.55 μm. The emission spectrum for the sky with stratus cloud resembles those of a black body and can be identified within the 8–14 μm spectral range of our infrared cameras. The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

has a smaller extinction coefficient [44]. In combination with their minor vertical extent, cirrus clouds are optically thin and appear almost transparent at infrared wavelengths. This is represented in Figure 2.11 by comparing the emission spectrum of a clear sky to a sky containing cirrus cloud. There are several contributing factors that all act to reduce the visibility of cirrus to a ground-based observer. Firstly, the emissivities of cirrus clouds are low (I calculate a value of 0.30 ± 0.12 in Section 6.1, which is in agreement with similar studies [28][45]). In addition, cirrus clouds are at high altitudes (10 km above ground level for this simulation) and consequently emit at cooler temperatures. As they are high in the atmosphere, there is also more intervening water vapour. In fact, the emission spectrum closely resembles that of a clear sky with water vapour (see Figure 2.9). It follows that cirrus clouds appear faint at the wavelengths observed by our cloud cameras. This is especially true when there is more humidity in the atmosphere. For the reasons already discussed, the difficulty in detecting cirrus clouds also increases with atmospheric path length.

The trouble with identifying optically thin clouds is similarly encountered in the analysis of infrared satellite data at Auger (see GOES, Section 3.3). In this case, thin clouds produce a negligible change in the overall infrared brightness radiated by the Earth's surface [46]. However, as noted in [46], thin clouds do not always have a recognisable effect on single extensive air shower events. Optically thin clouds, however, can instead be detected by the CLF and XLF (see Section 3.5) at the Pierre Auger Observatory.

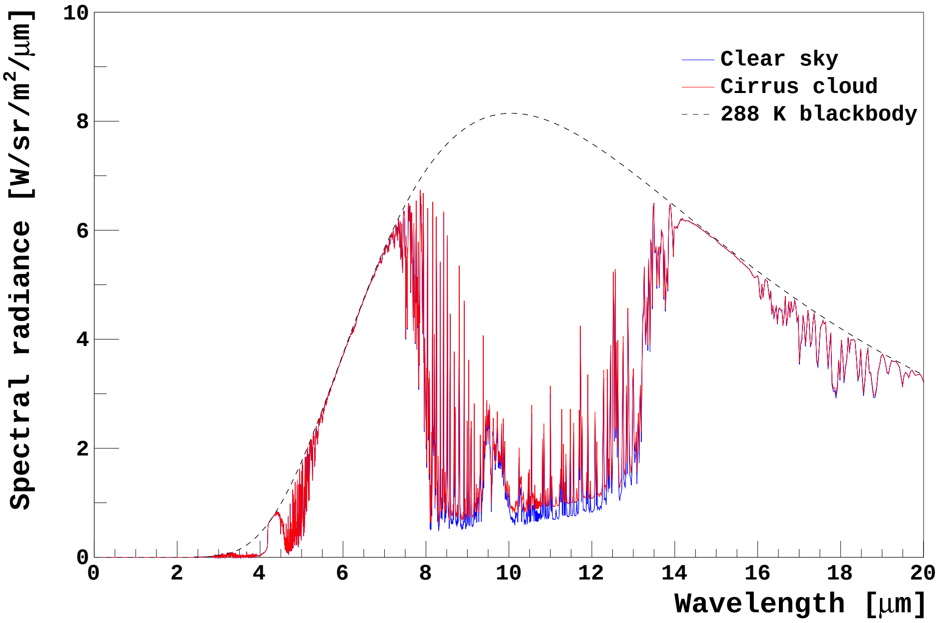

Figure 2.11: Atmospheric radiance reaching the ground from a clear sky, and a sky with cirrus cloud at 10 km above surface level. Simulated along a vertical track by the author with MODTRAN for a 1976 U.S. Standard Atmosphere. The cirrus cloud in this model has an optical depth of 0.14 at 0.55 µm. The emission spectrum for the sky with cirrus cloud resembles that of a clear sky with water vapour (see Figure 2.9). This leads to difficulties in their detection by our infrared cameras. The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

To demonstrate the effectiveness of our thermal imaging cameras at detecting clouds, Figure 2.12 shows some comparisons between clouds at (a) visible and (b) infrared wavelengths. Not only is there excellent agreement, but the infrared images are slightly more conservative, and show more cloud than the visible images. The images were taken in the evening and the morning, so that clouds were visible with the optical camera. Optical cameras are even less useful for detecting clouds at night-time. For the reasons outlined in this chapter, we can conclude that our infrared cameras are suitable for cloud detection at the Pierre Auger Observatory.

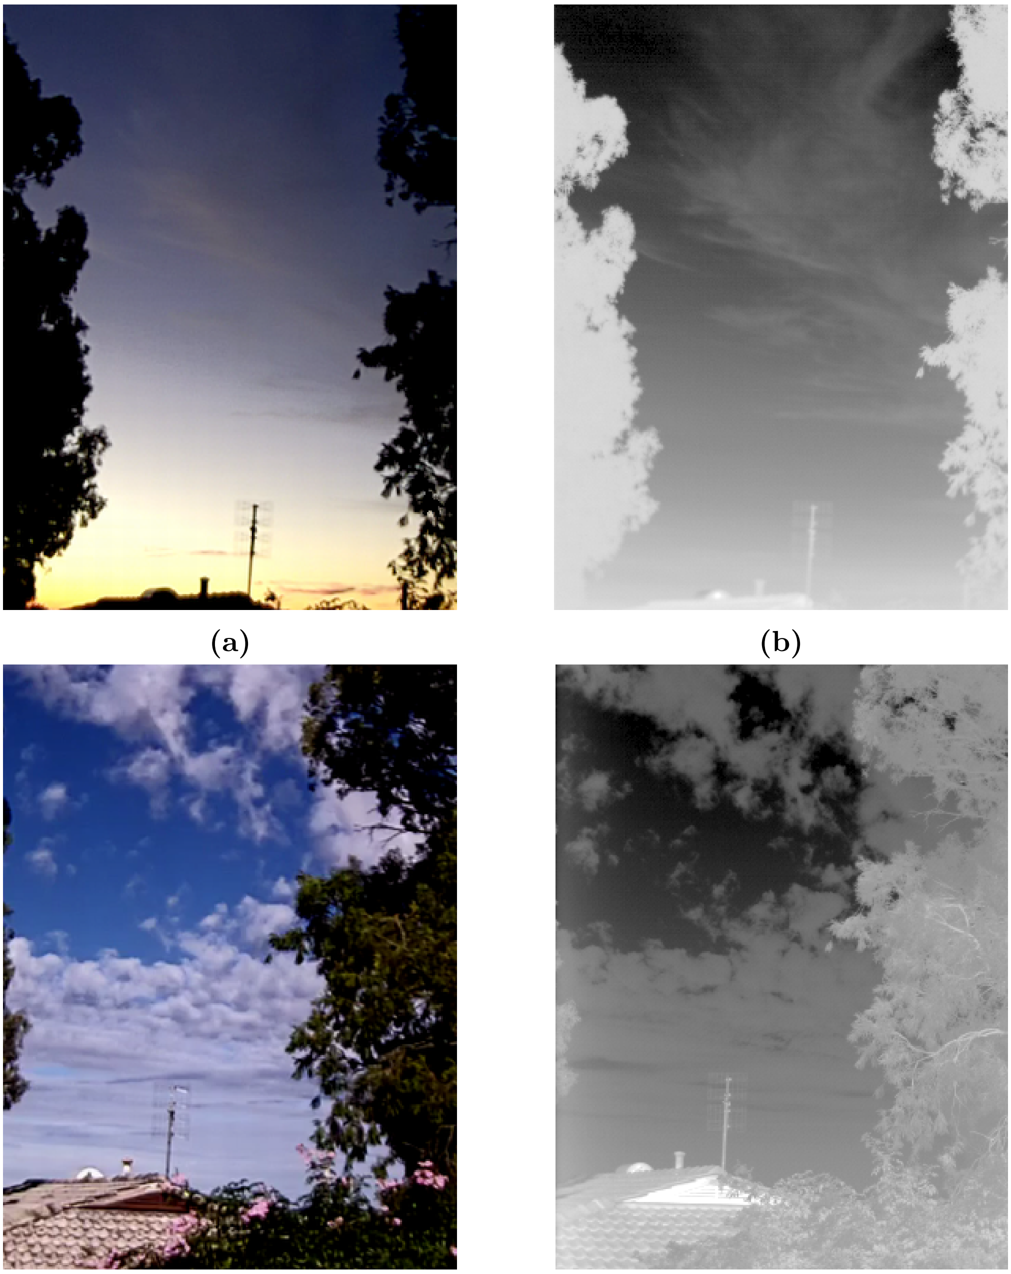

(a) **(b)**

(a) **(b)**

Figure 2.12: Comparisons between **(a)** visible and **(b)** infrared images of clouds. The top row of images were taken in the evening after sunset, and the bottom row of images were taken in the morning. The excellent agreement demonstrates that our infrared cameras are suitable for cloud detection. Images courtesy of Dr. Padric McGee.

Chapter 3

Meteorological Instruments used at the Pierre Auger Observatory

The atmosphere influences the development and detection of extensive air showers. This is particularly true for measurements made by the Auger fluorescence telescopes. The presence of clouds can lead to erroneous reconstruction of extensive air shower events (refer to Section 1.2.3). As the amount of observed fluorescence light relates to the amount of energy deposited into the atmosphere by an extensive air shower, it is important to know the atmospheric conditions. For this purpose, there is an extensive network of meteorological instruments used at the Observatory (see Figure 3.1).

Most relevant to my research have been the instruments capable of detecting clouds. Clouds may be passively observed through their thermal emission, such as with an infrared camera, radiometer, or weather satellite. They may also be detected from their scattering properties, such as with a lidar. The various instruments generally complement each other, and in practice can allow cross checks to be made.

3.1 Infrared cloud cameras

The basic principle behind an infrared detector is to convert incident radiation into some measurable signal. One way of doing this is to exploit some temperature-dependent property of a material. This idea can be extended to converting infrared radiance into quantitative images.

In the past, thermal imaging had been a costly endeavour. The devices

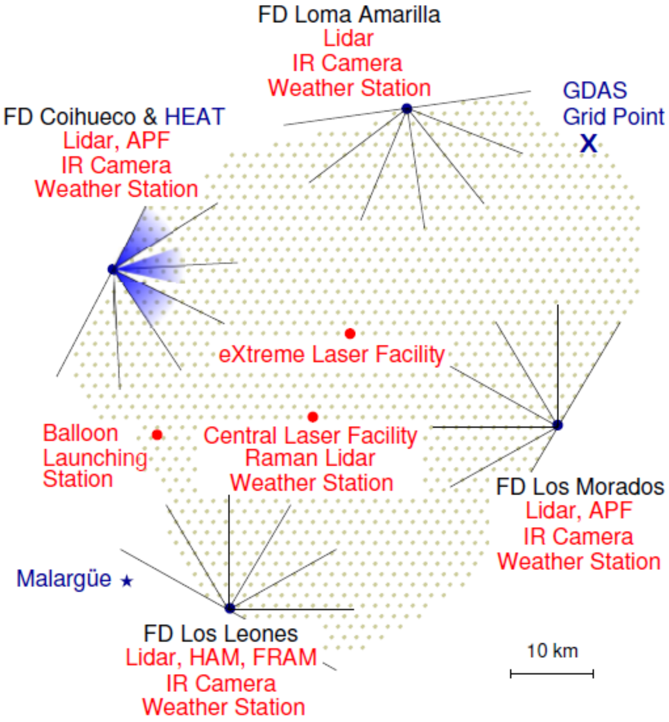

Figure 3.1: A schematic overview of the extensive network of meteorological instruments operating at the Pierre Auger Observatory. The following instruments for cloud detection are discussed in this chapter: the infrared (IR) cloud cameras, lidars, and the Central Laser Facility (CLF) and eXtreme Laser Facility (XLF) which both also contain a single-pixel infrared radiometer. Additionally, data from a weather satellite provides cloud cover data above the entire array. Image taken from [1].

were initially bulky, required external cooling, and were reserved for military applications. Over time, the field saw its way into commercial and industrial use. With modern advancements in microtechnology, the mass production of comparatively low-cost sensors has made thermal imaging affordable to a wider range of users. The compact design, relatively low maintenance and affordability of modern-day infrared cameras are suitable for our needs of night-time cloud detection at Auger.

3.1.1 Previous infrared cloud cameras

Clouds at long-wave infrared wavelengths are generally brighter than a clear sky (see Section 2.4). This means that clouds can often be resolved from a clear sky background. As part of the University of Adelaide's contribution to the Auger Observatory, infrared cameras have been installed at each of the main fluorescence detector (FD) sites to monitor the atmosphere for clouds. The first generation cloud cameras were the Raytheon 2000B series. The physical operation of those cameras were based on the pyroelectric effect (the temperature dependence of electrical polarisation). The pixels would be periodically exposed to different temperatures. This modulation of infrared radiation would be achieved with a mechanical chopper. The resultant heating-cooling cycle of a pixel would initiate a measurable flow of charge. Consequently, the signal would be proportional to the difference between the temperature in the field of view and the temperature of the chopper. Because there was no direct relation between the signal output and the temperature of a viewed object, these cameras were non-radiometric.

Although the cameras produced images capable of detecting clouds, there were several issues. The differential measurement, inherent in pyroelectric detectors, would lead to a difficulty interpreting an isothermal scene (such as a completely clear or overcast sky). Also relating to the design of the detector, the requirement for a chopper became problematic, as a misalignment in the chopper would manifest in the images. These resulted in undesirable image artefacts that could be recognised by eye, but were challenging for more automated methods of cloud detection. A detailed discussion on the Raytheon cameras and image analysis may be found in [22]. These cloud cameras operated until their replacement in 2013.

3.1.2 The Gobi-384 uncooled microbolometer array

The second generation cloud cameras for Auger are the Xenics Gobi-384 uncooled microbolometer array [47]. The model consists of a 384×288 pixel array that is sensitive to radiation within the 8–14 µm atmospheric window. It is capable of producing 16-bit images, and under normal operating conditions can typically resolve temperature differences of about 50 mK. We fit each camera with a 10 mm focal length lens to provide it with an approximate 54°×40° field of view. These new cameras are fundamentally different, as they operate using the bolometric effect instead of the pyroelectric effect. Unlike the old cameras, these are radiometric.

3.1.2.1 Physical concepts and pixel structure

Each pixel behaves as a (micro)bolometer. Incident radiation is absorbed and raises the temperature of the sensor. The electrical resistance of the bolometer is a temperature dependent property that can be measured. The active element of the sensor is a thin layer of amorphous silicon. To increase the spectral response near 10 µm, the sensor is suspended above a reflecting film. This separation acts as quarter wavelength cavity, such that any undetected flux is reflected back towards the detector [48]. The suspended bolometer membrane has a low thermal mass and is thermally isolated. This results in a low thermal time constant (a fast response to temperature variations) for the camera. A metal stud connects the sensor to a CMOS (complementary metal oxide semiconductor) readout integrated circuit (ROIC) [29]. These features are illustrated in Figure 3.2.

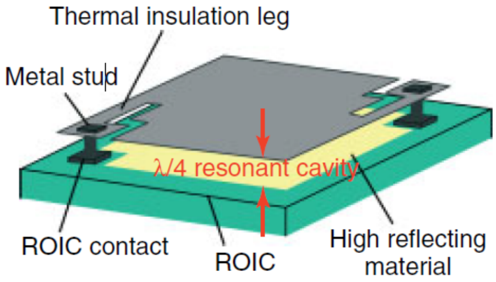

Figure 3.2: The structure of a microbolometer pixel [29], courtesy of ULIS.

The arrangement of individual microbolometers into columns and rows (384×288 pixels for our camera model) are coupled to the ROIC. This combined system forms the focal plane array (FPA). The ROIC reads out the pixel output row-by-row and amplifies the signal. To minimise detector noise, each active pixel is compared to a blind bolometer (one that is not illuminated) to subtract a common background current. To simplify the circuitry, a 12-bit analogue to digital converter (ADC) is integrated into the design of the ROIC. When writing the images, the 12-bit image resolution is scaled to have a 16-bit depth. The digital output, in ADC counts, relates to the infrared brightness within its field of view.

3.1.2.2 Camera settings and operating status

The ADC output of the camera depends on other factors besides the scene temperature. The settings of the camera, as well as the state of the camera, affect measurements. To analyse images in a quantitative manner, it is important to separate these effects from truly physical properties, such as the perceived temperature of an object, or atmospheric conditions.

For this purpose, the design of the ROIC includes a serial control bus. This provides the user with some degree of flexibility, by allowing for changes to be made to the settings through accompanying software, or for monitoring the status of the camera during its operation.

One quantity we have found to be important to monitor is the temperature of the FPA. This is measured by an on-chip thermometer and is representative of the temperature of the camera's sensor, which will be referred to as T_{cam} in this thesis. During normal operation, we find this typically stabilises to about 25 K above the environmental temperature (due to the camera's internal heater). Both the pixel output and temperature resolution depend on T_{cam} [49]. As a result, we find the response of the cameras to vary throughout a night, as the ambient temperature cools. We correct for the effects of the FPA temperature in the process of calibrating the cameras (see Section 4.1.2 and Section 4.5).

To optimise the response of our cameras to their environment, it was desirable to adjust some of the camera settings. This was largely a trial and error process, but has resulted in improving the quality of our images. We have found optimal settings by adjusting the image span, image level, and the voltage bias of the blind microbolometers (VSK). The span and level represent the image contrast and brightness, respectively [29]. To resolve smaller temperature differences in our 16-bit images, we maximise the contrast by setting the value of the span to 1 and the level to 65536. We found the VSK to be a sensitive parameter. To avoid image saturation, we now fix the VSK to a value of 3550 (units relating to voltage). The effects of all these quantities on the camera output were found experimentally. Apart from experimentation in 2013 and early 2014, the camera settings are now kept constant. Due to the effects these settings have on our cloud cameras, some care needs to be taken when comparing images from different epochs. A list of major changes made to the cloud camera software may be found in Appendix B.

3.1.2.3 Cloud monitoring system set-up and operation

The operation of the new cloud cameras is similar to the old system. A cloud camera is installed on the roof of each of the main fluorescence detector (FD) buildings, such that it can survey the FD field of view for clouds.

To protect the cameras from the environment, they are housed within a weather shield and view the sky through a H.VIR® crystal infrared window [50]. Consequently, the cameras also detect some thermal emission from the window material itself. The spectral response of the camera and the bandpass of the window are shown in Figure 3.3.

The housing is mounted on a pan-and-tilt system that drives the orientation of the camera. This set-up, shown in Figure 3.4a, allows the camera to

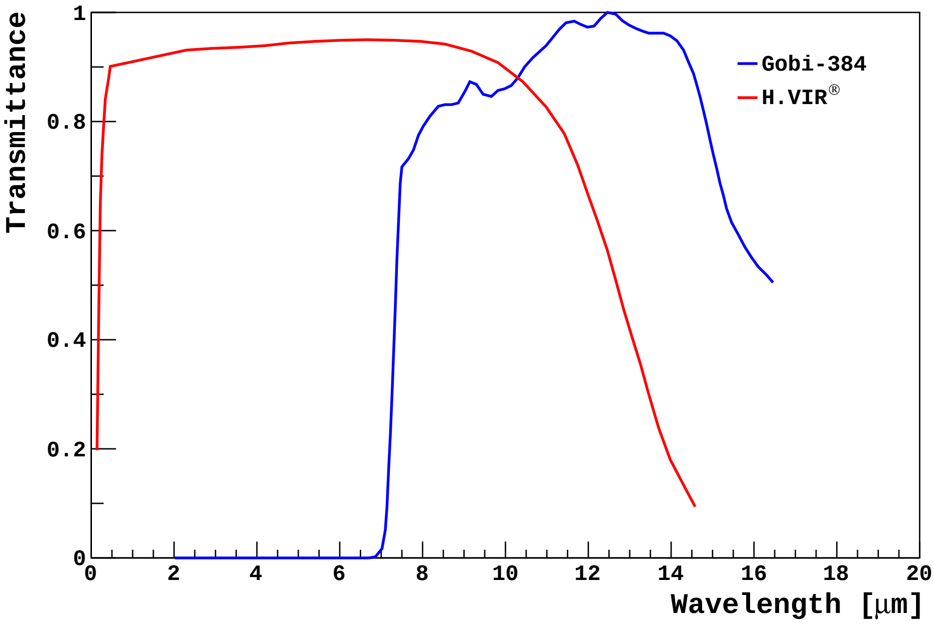

Figure 3.3: The spectral response of the Gobi-384 uncooled microbolometer array and the band-pass of the H.VIR® window. Data points for the camera and window have been extracted from [51] and [50], respectively.

view different regions of the sky. During FD data taking, the camera scans the FD field of view at 5 minute intervals. Given the 54° horizontal field of view of our cameras, we use 5 images to cover the entire FD field of view. This provides sufficient image overlap. Additionally, the cameras perform a 19 image full-sky scan every 15 minutes. This sequence consists of 10 horizontal images, 8 images elevated at 45°, and 1 image overhead. Figure 3.4b shows an example of a full-sky mosaic [1]. These provide information for the FD shifters on-site on the general cloud cover at the site, but are also useful for calibrating the cameras.

To operate the cloud monitoring units, each camera and pan-and-tilt are connected to a dedicated local computer within the local FD building and controlled by LabView software. We found the default software for the cameras to be limited for our needs of automated cloud monitoring. To routinely operate the cloud camera systems and provide detailed status reports, new software needed to be developed. This has been the work of Trent Grubb,

---

[1] Archived full-sky mosaics may be accessed from `http://www.physics.adelaide.edu.au/astrophysics/pierre/cloud.html`.

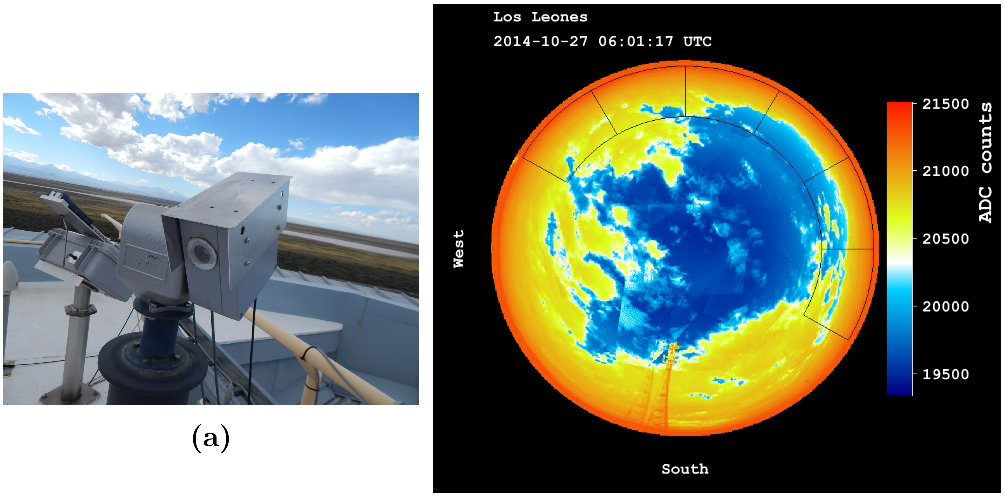

(a)

(b)

Figure 3.4: (a) The infrared cloud camera system at the Los Leones fluorescence detector site. The camera housing is mounted onto a pan-and-tilt unit driven by a motor, allowing the camera to view different regions of the sky. **(b)** An example of a full-sky mosaic at Los Leones, consisting of 19 individual images. The ADC counts relate to the infrared brightness in the camera's field of view. The black lines indicate the fields of view of the fluorescence telescopes.

a current PhD. student at the University of Adelaide. Another aspect of Trent's work has been to determine the pointing direction of each pixel for a given cloud camera image. This is based on tracking the Moon's position with a camera (passing through the field of view) and making comparisons to its actual position (which can be calculated). The pixel pointing directions (zenith angle and azimuth angle) are used in the subsequent work in this thesis.

The cloud camera software writes the images in the lossless PNG image format and saves them onto the local computer. The images are accompanied by an XML file that contains detailed information about a camera during the scanning process, including the time of the scan, the pointing direction of each image, and the camera settings. To be efficient in the amount of storage space we require, the files (PNGs and XML) for a given scan are then compressed. A typical night contains about 100 MB of data after this compression. The data are later transferred to Adelaide for future analysis, and to the Auger Computer Centre at Lyon.

3.2 Single-pixel infrared radiometers

In addition to our thermal imaging cameras, we have also installed two single-pixel infrared radiometers at the Pierre Auger Observatory. One of the radiometers is located at the Central Laser Facility (CLF), and another at the eXtreme Laser Facility (XLF). Our radiometers are sensitive to atmospheric thermal radiation, and can also be used to monitor the sky for cloud. We have previously seen that our cameras utilise a series of bolometers to achieve this effect. The radiometers used at Auger are instead based on the Seebeck effect.

3.2.1 Physical concepts and detector design

The Seebeck effect occurs when two dissimilar metals are in contact at two points [29]. One point (or junction) typically has an absorbing material that is sensitive to incoming radiation. The other junction is instead used as a reference. Electrons on the surface of one metal can be transported to the other. This requires energy, and the amount of energy depends on the material (i.e. the work function of the metal). This process upsets the charge balance of the metals and creates a contact potential. The number of electrons transferred depend on the temperature at the junction. If both junctions are at the same temperature, then the effects cancels out. However, when there is a temperature difference, the contact potentials are not equal. It follows that a temperature difference across the junctions can be converted into a potential difference.

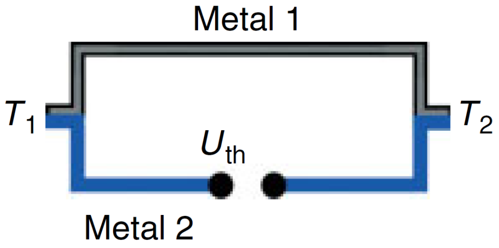

Figure 3.5: Two dissimilar metals, Metal 1 and Metal 2, make contact at two junctions. The temperatures at the junctions are T_1 and T_2. The Seebeck effect creates a measurable voltage U_{th} from a temperature difference across the junctions. This property is the basis for a thermocouple. Image taken from [29].

This property is the basis for a thermocouple (see Figure 3.5), as it can convert a temperature difference into a measurable voltage. Although the voltage produced is small, it can be amplified by stacking many thermocouples in series. This is known as a thermopile.

3.2.2 Instrument set-up at Auger

The University of Adelaide has developed simple infrared detectors based on the principal of the thermopile [52]. The focus of the design had been to inexpensively monitor for clouds at astronomical sites. For this purpose, the instruments use a single-pixel with sufficient sensitivity to distinguish cloud from a clear sky.

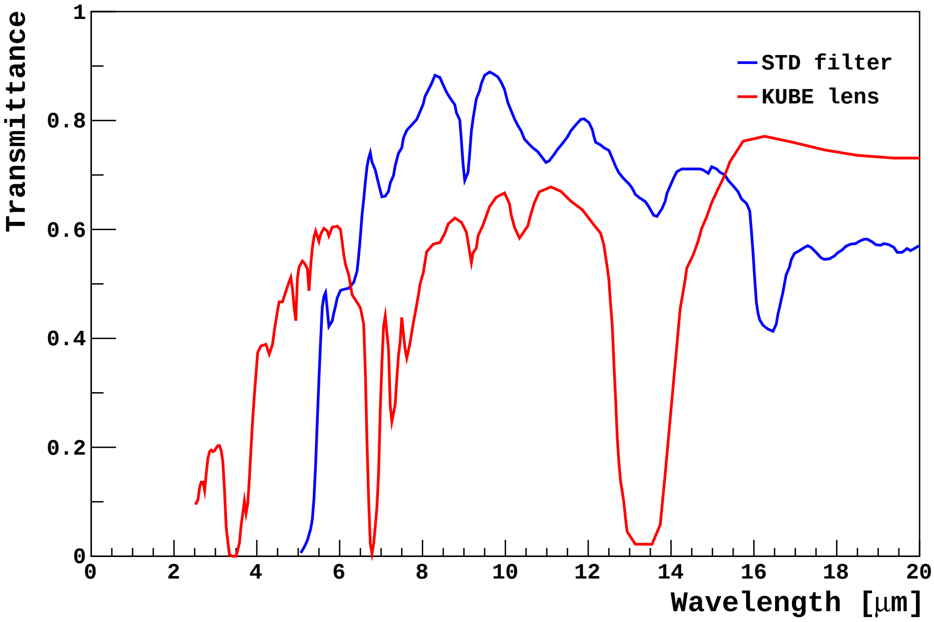

Figure 3.6: The spectral response of the TPS 534 thermopile with a standard (STD) filter, and the band-pass of the infrared fresnel lens (manufactured by KUBE). Data points for the detector and lens have been extracted from [53] and [54] respectively.

The active sensor element is the Perkin Elmer TPS 534 thermopile [55]. The detector has a sensitive area of 1.2 mm×1.2 mm, a response of 42 V W^{-1}, and is equipped with a standard filter that transmits all wavelengths above 5.5 µm. These cloud monitors are designed to operate both day and night in a variety of weather conditions. To protect the equipment, the sensor is housed in a die cast aluminium box. A thermistor is used to monitor the ambient temperature within the canister, which serves as the reference junction in the thermopile design. An infrared fresnel lens (manufactured by KUBE [54]) provides the instrument with a 3° field of view, and also shields the sensor from solar radiation, while still maintaining sensitivity at infrared wavelengths. The spectral response of the TPS 534 thermopile and the band-pass of the lens are presented in Figure 3.6. The somewhat similar spectral responses between our infrared cameras and the radiometers have been useful, particularly for calibrating the cloud cameras.

The detector maintains an output voltage proportional to the difference between the scene temperature and the thermistor temperature (the reference junction). This is known as the uncompensated output. Since the uncompensated output depends on the temperature difference between the thermocouple junctions, and the temperature of reference junction is known, the output can be compensated in order to determine the scene temperature. This is known as the compensated output, and corresponds to the temperature within in the detector's field of view. Each radiometer has been calibrated using an approximate black body, and the accuracy of the compensated temperature and thermistor value are about 1 °C [45].

The first single-pixel cloud monitor was installed at the Central Laser Facility (CLF) in 2005. The instrument is fixed near the ground, such that the internal temperature of the detector canister is approximately at ground temperature. The radiometer is directed towards the zenith in order to measure the infrared brightness of the sky, within its given spectral sensitivity. For a given cloud, the scene temperature depends on the cloud's height, its emissivity, and the temperature of the atmosphere at that height. Since the temperature of the atmosphere at that height will vary with ground temperature (assuming a constant lapse rate), the uncompensated output happens to be the most convenient output to use for cloud monitoring. A second instrument was similarly installed at the eXtreme Laser Facility (XLF) in 2008. Both radiometers record data at 5 minute intervals [2]. This fast response is particularly useful for cloud detection, as cloud conditions can vary quickly over a short period of time.

Low-level clouds have an effective sky temperature similar to the temperature of the ground. Consequently, an uncompensated temperature near 0 °C can indicate the presence of cloud. On the other hand, the uncompensated temperature for a clear sky is typically around −20 °C in our spectral range. Although intervening atmospheric water vapour increases this [53], a reasonable distinction can often be made between cloud and a clear sky background. Examples of radiometer measurements that are used to identify

---

[2] Data are available online from the Colorado School of Mines - Pierre Auger Laser Facilities: `http://astroserve.mines.edu/newindex.html`.

clouds are in Figure 3.7. We can infer that clouds were detected by the CLF radiometer on 2014-01-15 between approximately 00:40–01:40, 03:45–05:15, and 09:15–10:00 UTC due to the increase in uncompensated temperatures. In practice, the temporal information of clouds, atmospheric water vapour content [56], and an estimation of cloud base heights [45] can be inferred from the radiometers. My research involving the latter may be found in Section 6.1.

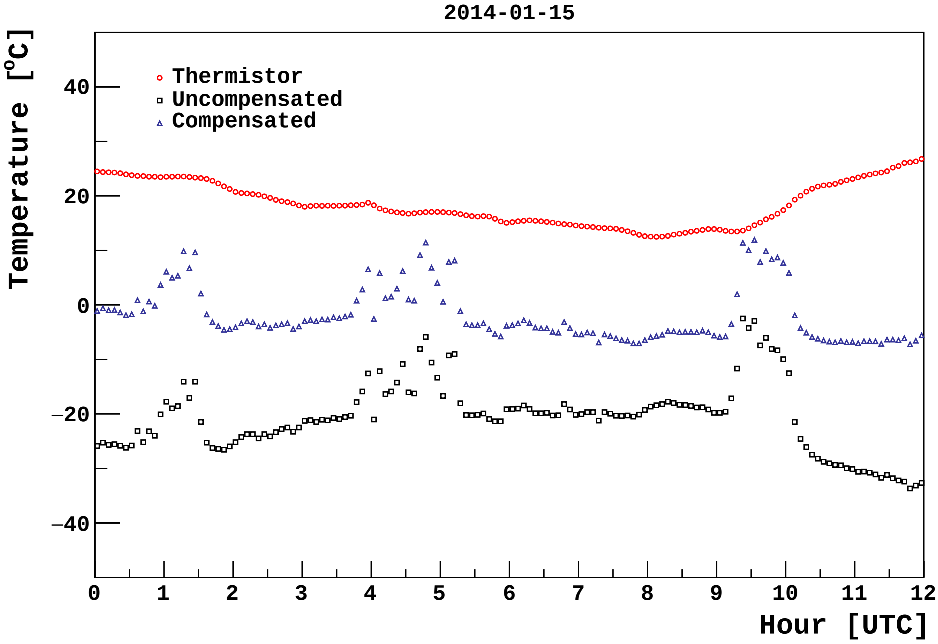

Figure 3.7: An example of cloud monitoring with the CLF radiometer on the 15th January, 2014. Variations in the uncompensated temperature can be used to infer that clouds were detected between approximately 00:40–01:40, 03:45–05:15, and 09:15–10:00 UTC.

3.3 Infrared satellite data

The detection of night-time cloud is also possible with infrared data from satellite instruments. At infrared wavelengths, the effective temperature of cloud is generally cooler than the ground. This is in contrast to the cloud cameras, that instead resolve cloud from being warmer than a clear sky.

A previous study has shown that raw satellite measurements made by the Geostationary Operational Environmental Satellites (GOES) can be used to detect clouds at the Pierre Auger Observatory [46]. Unlike the cloud cameras, information from GOES can be used to provide cloud cover above the entire observatory. These two instruments complement each other well.

3.3.1 GOES instrumentation

The GOES system is a series of satellites run by NOAA's (National Oceanic and Atmospheric Administration's) National Environmental Satellite, Data, and Information Service (NESDIS). For monitoring and forecasting the weather in the Americas, the program designates two satellites to simultaneously observe the Pacific Ocean and Atlantic Ocean. These satellites are referred to as GOES-West and GOES-East, respectively. We are particularly interested in data from the GOES-East satellite, as it captures the South American continent every 30 minutes. As of April 2010, the GOES-13 satellite has been operating as GOES-East. Note that at the time of writing my thesis, the GOES-16 satellite replaced GOES-13 in the GOES-East position on December 18th, 2017.

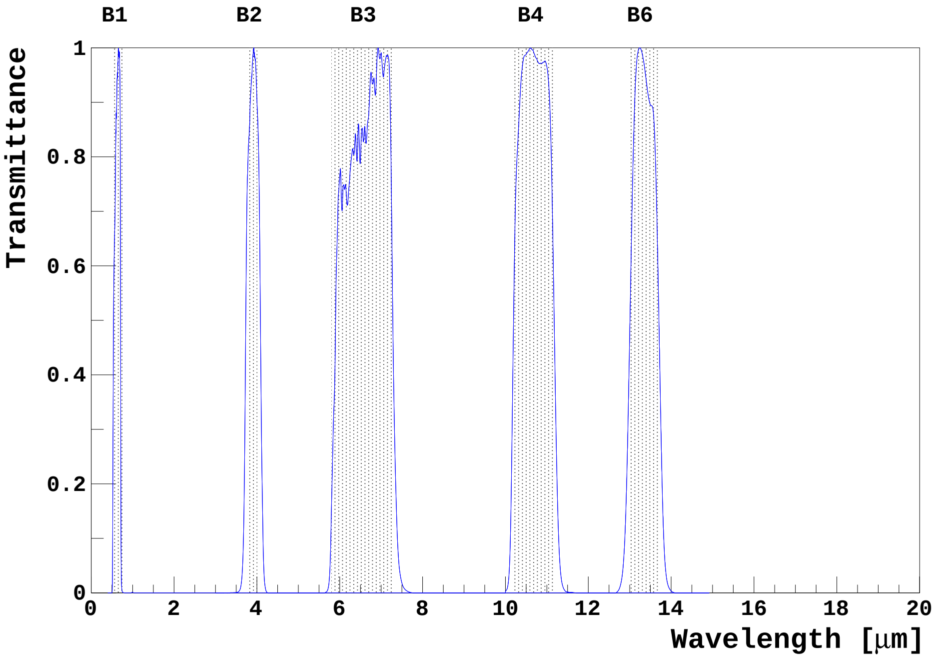

Figure 3.8: The spectral response of the GOES-13 Imager. One of the bands, B1, detects visible light. The four remaining bands, B2–B6 are sensitive to infrared radiation. Data points are taken from [57].

One of the instruments on-board the satellites is the multichannel imaging radiometer, Imager. It senses radiant and solar reflected energy from the Earth in five distinct bands (one visible and four in the infrared). For the GOES-13 satellite, the central wavelength of each band is at 0.65, 3.90, 6.55, 10.70, and 13.35 µm. These correspond to the visible (B1), short-wave infrared (B2), water vapour (B3), long-wave infrared (B4) and carbon dioxide (B6) channels [3]. The nadir spatial resolution of B6 is 8 km×8 km, while the remaining three infrared bands have a higher resolution of 4 km×4 km. The visible channel has a finer resolution of 1 km×1 km. The spectral response of the GOES-13 Imager can be seen in Figure 3.8.

3.3.2 Cloud detection at Auger using infrared GOES data

Most relevant for night-time cloud detection are the Imager instrument's infrared channels. At infrared wavelengths the Earth's surface radiates as a black body. On the other hand, the emissivity and temperature of cloud is typically lower, and the emissivity shows a stronger wavelength dependence. This dependence results from the size distribution of the cloud particles [43]. Consequently, the flux received by the instrument will vary across each of the channels in the presence of cloud. The temperature of the atmosphere at the height of the cloud is also typically cooler than the ground. If the cloud is optically thick, then radiation emitted from the Earth's surface is obscured, further reducing the flux (proportional to the fourth power of the temperature difference between the cloud and ground). This enhances the separation between a clear and cloudy sky. Figure 3.9 shows a model of the flux received in each of Imager's channels. It is often useful to convert the spectral radiance into a brightness temperature (see Equation 2.5 in Section 2.1), as this quantity is affected by both the temperature and emissivity of the emitting object. The corresponding plot in brightness temperature is shown in Figure 3.10.

It is common to use a combination of channels for the purpose of cloud detection. The non-absorbing bands, B2 and B4 are particularly useful, as they are the least affected by atmospheric attenuation. One method of detecting water clouds (and fog) at night is to measure the difference between the brightness temperature in the short-wave infrared band (T2 of B2) and the long-wave infrared band (T4 of B4) [59]. This subtraction is sometimes referred to as the fog/stratus product. This technique is sensitive to emissivity differences between the bands, as they respond to temperature in the

---

[3] Prior to GOES-12, a 12 µm channel (B5) was used instead of the carbon dioxide channel [58].

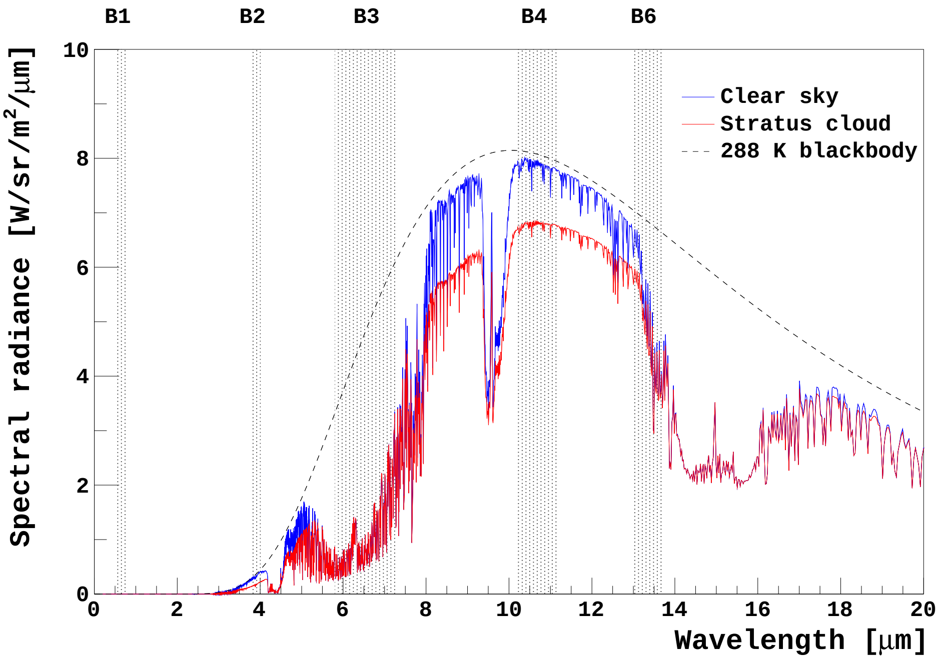

Figure 3.9: The radiance of the Earth as seen by a geostationary satellite. Simulated by the author using MODTRAN with a 1976 U.S. Standard Atmosphere and stratus with a cloud base at 0.33 km and a cloud top at 1 km. The shaded regions are observed by the GOES-13 Imager instrument. The stratus cloud in this model has an optical depth of $\sim$ 38 at 0.55 µm. The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

same way. Since the emissivity of cloud is not constant with wavelength, it follows that in the presence of cloud, T2 and T4 will not be the same. Temperature differences T4-T2, associated with stratus cloud and fog are typically 2–5 K [59]. On the other hand, for a clear sky, this temperature difference is caused by differential water vapour absorption. This is a smaller effect. Both cases, for stratus cloud and a clear sky with water vapour, are shown in Figure 3.10.

The Earth's outgoing radiation is highly attenuated by water vapour around 6.5 µm (discussed in Section 2.2.1). As the composition of clouds is a mixture of water vapour and liquid water droplets, they can be detected by Imager's water vapour channel, B3. This is also sensitive to atmospheric moisture, not just within clouds. As a result, B3 typically does not penetrate deep into the atmosphere. In the presence of water vapour, the atmosphere becomes optically thick and radiates as a black body at a temperature cor-

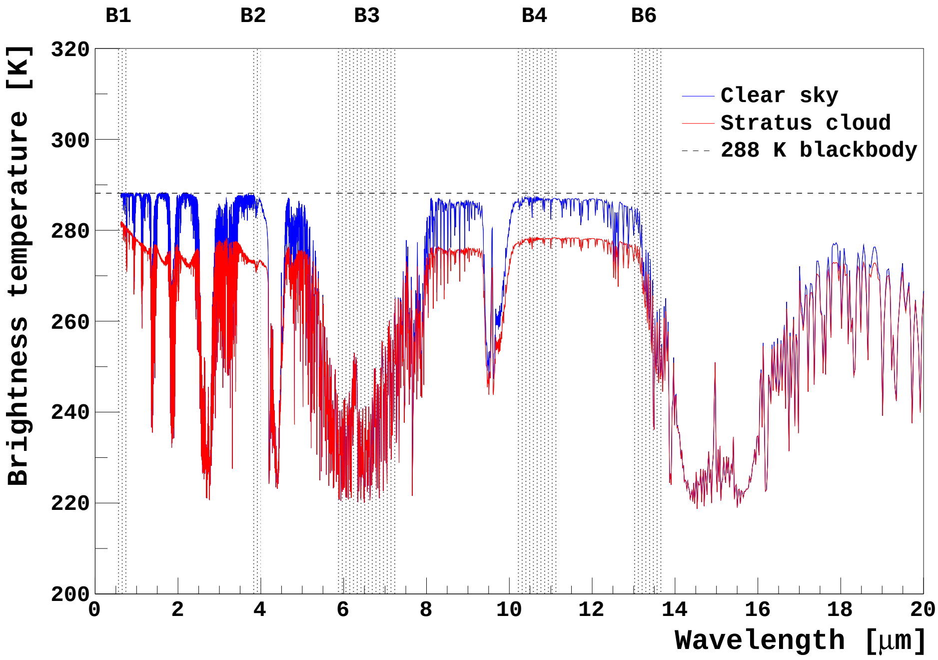

Figure 3.10: The brightness temperature for the same atmosphere as Figure 3.9. There is a temperature difference of ∼ 5 K between 10.7 µm (B4) and 3.9 µm (B2) for stratus cloud. Since the cloud is optically thick, the sensor measures thermal emission from the cloud top. The temperature of the atmosphere at the cloud top is about 280 K in the model. The low-level cloud (stratus) has the same signal as a clear sky at 6.5 µm (B3), as any outgoing radiation is significantly attenuated by atmospheric water vapour.

responding to its height in the atmosphere. This can be seen in Figure 3.10 from the significantly lower brightness temperature within the water vapour channel (T3 of B3). This channel instead probes the upper troposphere and is most suited for detecting high-level cloud.

A study investigating cloud detection at Auger using data from GOES has found that a combination of brightness temperatures, T2-T4 and T3 provide good discrimination between cloudy and clear satellite pixels [46]. In that analysis 2007 data were used from the GOES-12 Imager instrument [4]. Each satellite pixel, with the associated brightness temperature in each channel, was mapped onto the Pierre Auger Observatory with a spatial resolution of 2.4 km×5.5 km.

To relate this set of brightness temperatures to either cloudy or clear

---

[4] The GOES-12 and GOES-13 Imager have the same spectral response.

conditions, the authors used the Central Laser Facility (CLF) to provide a ground-truth [60]. The CLF operates in tandem with the Auger fluorescence detectors (FDs) to identify cloud within a FD's field of view (see Section 3.5). Of particular interest were observations of cloud directly above the CLF, as it matches the perspective of the satellite observations. With this, the cloudy or clear state of the satellite pixel encompassing the CLF could then be appropriately selected.

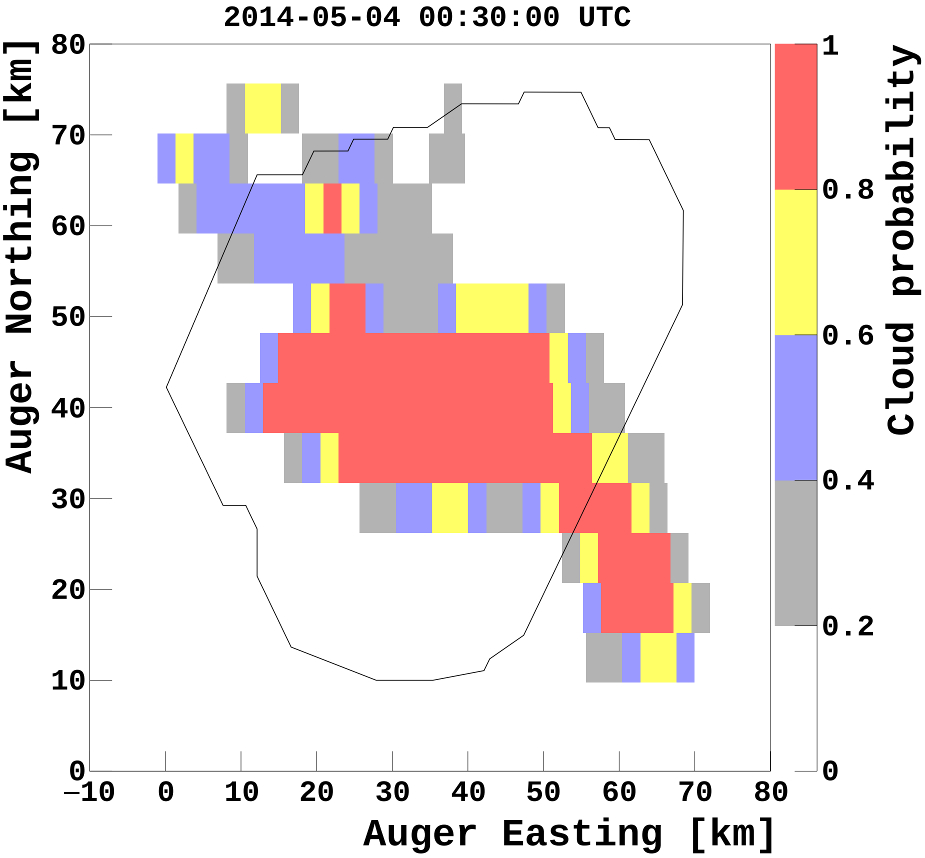

Figure 3.11: An example of a cloud probability map. The outline of the Pierre Auger Observatory is in black. A bank of cloud is seen in the middle of the array. Illustrated by the author of this dissertation.

Following this identification, the authors obtained separate distributions for cloudy and clear measurements that are dependent on their T2-T4 and T3 brightness temperatures. These represent probability distribution functions, and give the likelihood of the pixel containing cloud. This method of satellite-based identification can be extended to provide cloud probabilities for the remaining satellite pixels that cover the Auger Observatory. For this purpose a 30×12 array of pixels is used to provide cloud cover above the entire site. Figure 3.11 shows one example of a cloud probability map.

For further cosmic ray analysis, the cloudiness of each GOES pixel, at a given time, is stored in an Auger $\overline{\text{Off}}\underline{\text{line}}$ atmospheric database. In-depth reports on constructing these cloud probability maps can be found in [46][60].

3.4 Elastic backscatter lidars

The cloud cameras measure the angular position of cloud, but do not directly provide information on the cloud height. One instrument that can determine the height of cloud is a lidar. Unlike the previously discussed detectors that sense thermal emission from clouds, a lidar instead uses a laser beam to actively measure clouds from their scattering characteristics. These two instruments complement each other, and together can provide useful geometry for reconstructing cosmic ray events.

3.4.1 The Auger lidar system

The Auger lidar system consists of four elastic backscatter lidars, with one installed at each of the main fluorescence detector (FD) sites [61]. Each of the lidars uses a pulsed 351 nm laser beam to measure the attenuation and scattering properties of the atmosphere within the spectral range of the fluorescence telescopes. The extent of these effects vary in the presence of clouds and aerosols, and may be found by analysing the backscattered signal.

To collect this backscattered light, each lidar system contains a set of three parabolic mirrors. Each mirror has an associated photomultiplier onto which it focuses the light. As the transmitter and receiver are co-located, each lidar is monostatic (to distinguish the technique from the bistatic CLF and XLF lidars, discussed in Section 3.5). To measure the properties of the atmosphere surrounding the FD site, the lidar station is mounted on a steerable frame. This allows the lidar to perform a series of automated scans during FD data acquisition. The hourly routines have both a continuous and discrete scanning mode. For the purpose of detecting clouds, the continuous scan sweeps across the sky along two fixed orthogonal paths. To avoid contaminating measurements made by the fluorescence detectors, the maximum zenith angle along both paths does not exceed 45° (see Figure 3.12). Following the continuous scan is a series of discrete scans. In this mode the lidar telescope is directed towards different coordinates in order to accumulate larger statistics used for aerosol analysis.

3.4.2 Method of cloud detection

The scanning routine allows for hourly measurements of the cloud cover, cloud heights, and optical depths of the observed cloud layers. These quantities may be found by identifying clouds in the backscattered signal. Clouds are characterised by the lidar as strong localised sources of scattering, and can typically be detected. The cloud heights can be calculated from timing information of the returning signal. One example showing this is Figure 3.12, where the cloud base can be observed to be $\sim 3.5\,\text{km}$ above the lidar.

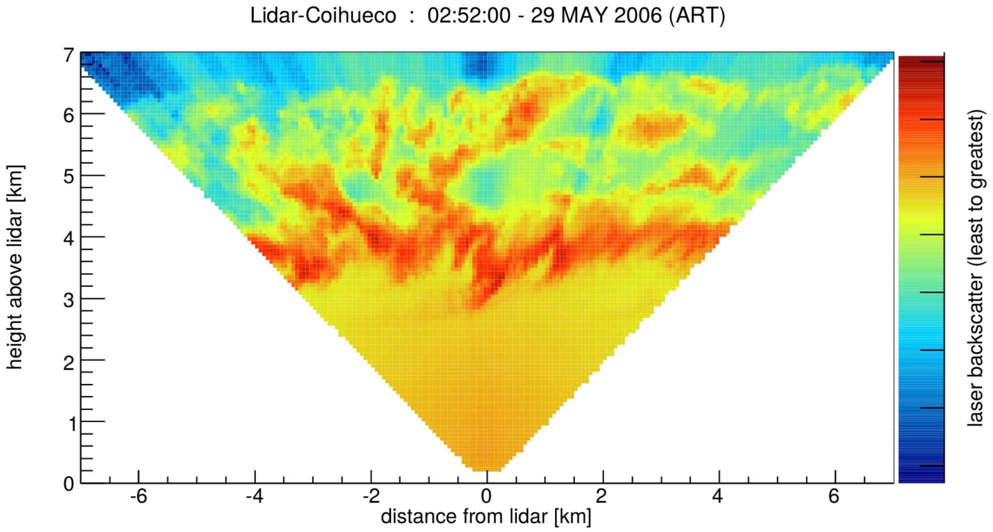

Figure 3.12: Cloud detected during a continuous scan made by the lidar at the Coihueco fluorescence detector site. The cloud layer is characterised as a strong source of backscattered light. Image taken from [61].

One method of automating the cloud detection is to compare the reflected signal to one expected for a clear sky. This is achieved by simulating the backscattered light for a purely molecular atmosphere. The scattering properties for a molecular atmosphere are empirically derived from monthly models made by radiosonde launches at the Pierre Auger Observatory [24]. The residual signal, found by subtracting the clear sky response from the observation, can be identified as cloud. Sudden variations in the residual signal can be used to locate the height and thickness of cloud layers in the atmosphere [62]. To reduce signal noise and limit the possibility of false detections, clouds must be detected by more than one lidar mirror.

Following the cloud identification, the lowest cloud layer is determined and the fractional cloud cover, based on the scan, is calculated. The signal difference between the base and the top of the cloud layer relates to its transmittance, which is used to estimate the mean optical depth within the cloud.

Each lidar station reports the hourly minimum cloud base height, the cloud cover, and the cloud optical depth. If no clouds were detected, the maximum range reached by the lidar is used to place a lower limit on the cloud base height. This information is filled into an Auger $\overline{\mathrm{Off}}\underline{\mathrm{line}}$ database, so that it may be used during the reconstruction of cosmic ray showers. Further details regarding the design and operation of the Auger lidar system may be found in [61].

3.5 Central Laser Facility and eXtreme Laser Facility

The Auger lidars operate outside of the fluorescence detector (FD) fields of view to monitor the sky for cloud. For cosmic ray analysis, there is particular interest in clouds that can potentially affect measurements made by the FDs. Cloud, as seen by the fluorescence telescopes, can be detected with the infrared cloud cameras. This is complemented with cloud identification using a bistatic lidar technique, where transmitted light from the Central Laser Facility (CLF) and eXtreme Laser Facility (XLF) is received at each of the FD sites.

The CLF and XLF are centrally located in the surface detector array. Both facilities generate a pulsed 355 nm laser beam that serves as a ”test beam” for a number of monitoring tasks [63]. As the laser light is attenuated in the same way as fluorescence light from an air shower, it can be used to determine the optical properties of the atmosphere, within the spectral sensitivity of the detectors. For the purpose of atmospheric monitoring, the CLF and XLF fire sets of 50 vertical laser shots into the sky every 15 minutes during FD operations. Due to the molecular and aerosol composition of the atmosphere, a small fraction of light scatters out of the laser beam and towards the FD sites, illuminating some of the telescope pixels. The triggered pixels form a track across the detectors. The amount of light recorded along the laser track, combined with timing information, is used to create light profiles (some examples are shown in Figure 3.13). The shape of these profiles provides key information in determining the scattering properties of the atmosphere.

The overall shape of the light profiles is governed by interactions between photons that have left the beam, and the intervening atmosphere on their paths towards the detectors. It follows that light arriving later at a detector must travel a greater distance through the atmosphere, and consequently, suffers greater attenuation. As a result, the intensity of the light received decreases smoothly with time (see Figure 3.13a). The extinction of light, at the given wavelength of the laser, is predominantly due to scattering off air molecules and aerosols. The former is well understood. However, the distribution of aerosols in the atmosphere varies over short time-scales, and their scattering process cannot be described analytically. Determining the scattering contributions due to the aerosols and their vertical distribution above the array requires frequent measurements. This is a fundamental goal of the CLF and XLF. Another goal is the identification of cloud.

3.5.1 Methods of cloud detection

Cloud can be detected as deformations in the otherwise smoothly falling light profiles. Two situations can arise, depending on the position of the cloud, relative to the laser beam and the detector. Photons entering a layer of cloud directly above the beam will undergo multiple scattering. This enhances the amount of light scattered toward the detector, and results in a characteristic peak that can be observed in the light profile (see Figure 3.13b). In this case the location of the cloud is known, so there is sufficient information to determine the height of the cloud. On the other hand, cloud positioned somewhere between the beam and the detector will obscure the light. The reduced intensity produces a dip in the light profile (see Figure 3.13c). Unlike the former scenario, the location of the cloud is not known. This ambiguous geometry means that the cloud height cannot be determined with this method. An upper limit of the cloud height, however, can be placed. To detect these cloud features, and to calculate the vertical aerosol optical depth (VAOD), two independent techniques are used: the Data Normalised method and the Laser Simulation method.

The approach taken in the Data Normalised method is to compare hourly averages of the light profiles to a reference [64]. This reference is constructed from profiles recorded on clear, and aerosol free, nights. In this procedure, each laser shot is normalised to an energy of 1 mJ. The measured laser profiles, as a function of height (in 50 m height bins), can be determined from the event timing and the known positions of the FD site and laser. For a given hour, each set of 50 laser shots is grouped together to create four, quarter hour profiles. Clouds are then identified in each height bin by comparing the photon transmission in the quarter hour $T_{quarter}$ to that of the reference T_{ref}. A ratio of $T_{quarter}/T_{ref}$ less than 0.1 indicates a dip in the profile, suggesting that light has been lost due to cloud positioned somewhere

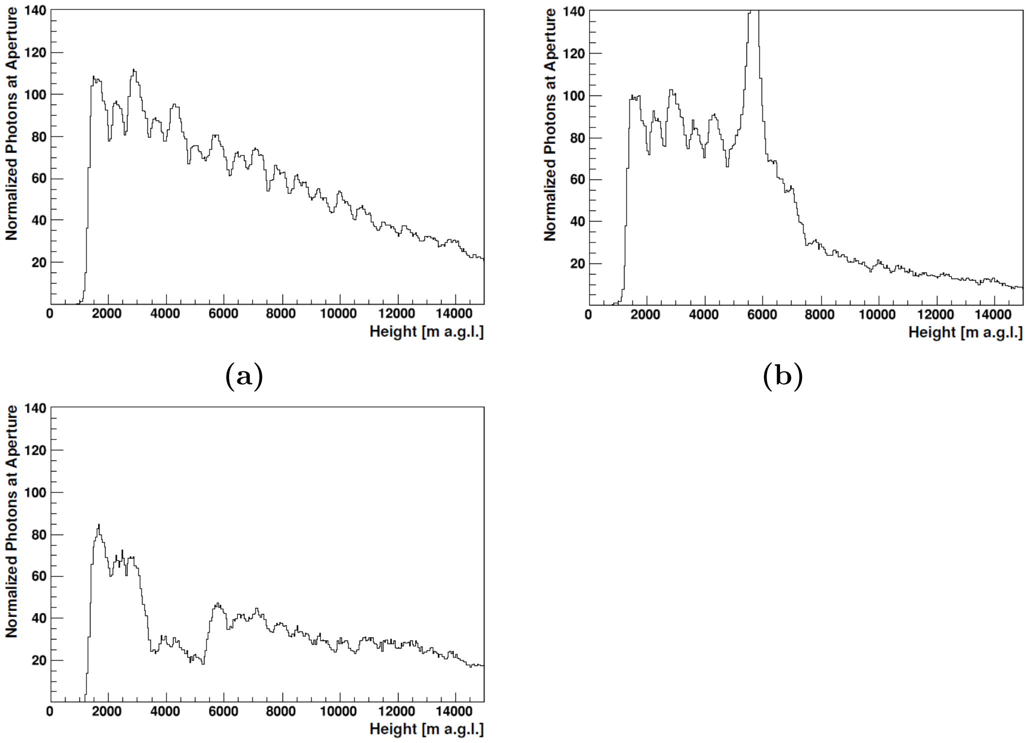

(a)

(b)

(c)

Figure 3.13: Examples of Central Laser Facility (CLF) light profiles, as seen from the Coihueco fluorescence detector (FD) site. **(a)** Clear sky with a smoothly falling profile. **(b)** Cloud directly above the CLF, showing a peak in the profile. **(c)** Cloud positioned somewhere between the CLF and FD site, producing a dip in the profile. These features are explained in the text. The small modulations in the light profiles are caused by the light collectors surrounding the PMTs. Images from [64].

between the laser and the FD. On the other hand, a ratio of $T_{quarter}/T_{ref}$ greater than 1.3 indicates the laser beam directly entered a layer of cloud. The algorithm does not currently distinguish between both cases. In either event, the minimum cloud base height is reported to be at the beginning of the anomaly.

The principle behind the Laser Simulation method is to instead compare the quarter hour profiles to simulations [64]. In this procedure, the model that best describes the atmosphere is determined from a series of simulations generated under different aerosol attenuation conditions [65]. Cloud can subsequently be identified by analysing the difference between the measured light profile and the best-fitting simulated profile. Much like the Data

Normalised analysis, peaks and dips are used to infer the presence of cloud. The minimum cloud base height is determined from the height bin containing the highest (or lowest) signal and the greatest signal to noise ratio.

Both the Data Normalised method and Laser Simulation method combine the quarter hour sets into a full hour profile. If clouds are identified in at least two of the quarter hours, then the full hour is regarded as cloudy. In that case, the minimum cloud base height found during the hour is reported. As we have discussed, there is some ambiguity in the cloud height for cloud associated with dips in the light profiles. In that situation it is important to note that the reported values do not necessarily represent their true height. If no cloud was identified in the hour, then a lower limit on the cloud height is placed at the height corresponding to the top of the FD's field of view. This hourly information is stored into an Auger $\overline{\text{Off}}\underline{\text{line}}$ database for further cosmic ray analysis.

3.6 The Global Data Assimilation System

We have primarily discussed the impact that clouds can have on the development and detection of extensive air showers (see Section 1.2.3). However, atmospheric state variables such as temperature, pressure, and humidity also affect measurements made with the fluorescence technique [24].

For a meaningful reconstruction of cosmic ray events, measurements made by the Auger fluorescence detectors (FDs) require the correct treatment for these atmospheric effects. In addition to the ground-based instruments and the weather satellite used for atmospheric monitoring, height-dependent profiles of these atmospheric state variables are also needed.

In the past these atmospheric profiles were obtained from radiosonde launches done by the collaboration at the Auger site. The data collected from an extensive weather balloon program were averaged to produce local monthly models [24]. Although these monthly models provided a broad seasonal description of the atmosphere, they could not account for day-to-day variations, which can be significant [66].

A previous study at the Auger Observatory has found that height-dependent atmospheric profiles obtained from the Global Data Assimilation System (GDAS), a global atmospheric model, were a suitable replacement for the monthly models [21]. The GDAS model offers a finer time resolution than the monthly models, with new meteorological data available every 3 hours. Applying this GDAS data to air shower reconstructions has been found to substantially reduce systematic errors and overall uncertainties, when compared to using the monthly mean profiles [21].

3.6.1 Impact of atmospheric state variables on measurements

Major atmospheric state variables such as temperature, pressure, and humidity, influence the development and, more importantly, the detection of extensive air showers [24][67].

The Auger FDs are designed to measure fluorescence light emitted from atmospheric nitrogen excited by the passage of air shower particles. This process of radiative de-excitation is quenched by collisional de-excitation of nitrogen with other atmospheric molecules, such as nitrogen, oxygen, and water vapour. These interactions are dependent on atmospheric temperature, pressure, and water vapour content. As a result, the fluorescence yield is sensitive to these state variables [20][68].

Moreover, depending on the geometry of the extensive air shower, measurements made by the FDs may be contaminated by Cherenkov light emitted by charged secondary air shower particles passing through the atmosphere. If the number of Cherenkov photons is known, this can be compensated for [69]. The Cherenkov yield depends on the refractive index of air, which also depends on temperature, pressure, and humidity of the atmosphere [70]. A simple parametrisation for the total refractive index of atmospheric air (at a particular density) n_{tot} in terms of dry (and CO_2-free) air, carbon dioxide, and water vapour content is given by:

$$n_{tot} - 1 = (n_{dry} - 1) \cdot \frac{\rho_{dry}}{\rho_{air}} + (n_{CO_2} - 1) \cdot \frac{\rho_{CO_2}}{\rho_{air}} + (n_w - 1) \cdot \frac{\rho_w}{\rho_{air}}\,, \tag{3.1}$$

where the refractive index of each component is weighted by its density. The number density of the atmospheric constituents decrease with height, which can be parametrised in terms of temperature and pressure. This can be used to determine the refractive index of air n_{air} [71].

In both processes, the transmission of the fluorescence and Cherenkov light towards the FD telescopes suffers from atmospheric attenuation. At ultraviolet wavelengths this is largely governed by Rayleigh scattering with atmospheric molecules. The Rayleigh cross-section σ_R depends on atmospheric temperature T, pressure P, and water vapour pressure e [72], and can be expressed as:

$$\sigma_R(\lambda, P, T, e) = \frac{24\pi^3}{\lambda^4 N^2} \cdot \left(\frac{n_{air}^2 - 1}{n_{air}^2 + 2}\right)^2 \cdot F_{air}(\lambda, P, e)\,, \tag{3.2}$$

where λ is the wavelength of the light, N the atmospheric molecular density, and F_{air} the King correction factor used to account for the anisotropies in the air molecules.

Apart from impacting on the detection, the state of the atmosphere also influences the air shower development. The interactions and decays of the secondary particles depend on the atmospheric depth, the total column density of atmospheric matter traversed. The atmospheric depth X can be calculated by integrating an air density profile $\rho(h)$ from ground level h_0 along a track. This is given by:

$$X(h_0) = \frac{1}{\cos\theta} \int_{h_0}^{\infty} \rho(h) dh \,, \tag{3.3}$$

where the zenith angle θ relates to the trajectory through the atmosphere. To the extent that the atmosphere behaves as an ideal gas, the air density has a dependence on temperature, pressure, and the molar mass of air, which may include water vapour.

Atmospheric conditions also impact on measurements made by the Auger surface detectors. An increase of pressure at ground level corresponds to an increased amount of matter traversed by the air shower particles. This affects the development of the air shower. On the other hand, variations of the air density affect the Molière radius. The latter determines the lateral spread of the extensive shower. Consequently, surface detector event rates are modulated by pressure and air density [73]. To correct for these effects, weather station temperature and pressure data can be used. Although there is a relationship between air density (at a given height) and surface temperature (measured by a ground-based weather station), the amplitude of the correlation depends on the height in the atmosphere. In addition, there is a delay in response to ground temperature variations [67]. The extent of the modulation can be understood from vertical temperature profiles.

In summary, the effects these state variables have on the development and detection of extensive air showers can be appropriately quantified with height-dependent atmospheric profiles. One approach is to use meteorological data from the GDAS model.

3.6.2 The GDAS model

The Global Data Assimilation System (GDAS) is an atmospheric model developed at NOAA's (National Oceanic and Atmospheric Administration's) National Centers for Environmental Prediction (NCEP). The system is an iterative process used to incorporate real observations into the Global Forecast System (GFS), a numerical weather prediction model.

The first step involves collecting meteorological data from a variety of sources around the globe. This includes data from weather stations, radiosondes, aircraft reports, weather buoys, and weather satellites (such as the

GOES satellites) [74]. The current observations, combined with additional information from the previous iteration of the numerical weather prediction model, are used as a first guess to forecast atmospheric conditions for a future state. The model forecast is adjusted by new observations to provide a better description of the atmosphere. This new model serves as the first guess in the next iteration.

This assimilation process of combining measurements with forecasts produces a 3-dimensional representation of the atmosphere. The model describes the atmospheric conditions for a given time, at a set location. The data are available every 3 hours, globally on a $1°\times1°$ latitude-longitude grid. Each data set contains both surface (S) and upper level (U) values of atmospheric parameters. The meteorological fields are listed in Table 3.2 (at the end of this section). The upper level profiles are interpolated from the GDAS model at 23 fixed pressure levels, which extend from 1000 hPa to 20 hPa. The pressure levels and their approximate heights (a.s.l.) are shown in Table 3.1.

Level	**Pressure [hPa]**	**Height [km]**	**Level**	**Pressure [hPa]**	**Height [km]**
23	20	26.4	11	600	4.2
22	50	20.6	10	650	3.6
21	100	16.2	9	700	3.0
20	150	13.6	8	750	2.5
19	200	11.8	7	800	1.9
18	250	10.4	6	850	1.5
17	300	9.2	5	900	1.0
16	350	8.2	4	925	0.8
15	400	7.2	3	950	0.5
14	450	6.3	2	975	0.3
13	500	5.6	1	1000	0.1
12	550	4.9	0	surface	

Table 3.1: The 23 constant pressure levels in the GDAS data. For reference, the corresponding heights of the 1976 U.S. Standard Atmosphere are also listed in this table. Pressure level 0 contains the surface values. Table adapted from [21].

For the Auger site, the most suitable GDAS grid point is positioned northeast of the array at (35°S,69°W), as shown in Figure 3.14. From the consistency shown between the various weather stations, it can be inferred that the atmosphere is horizontally uniform across the array [24][75]. Therefore, the atmospheric conditions at the selected grid point can be extended to the entire Observatory. The quantities most relevant for cosmic ray analysis

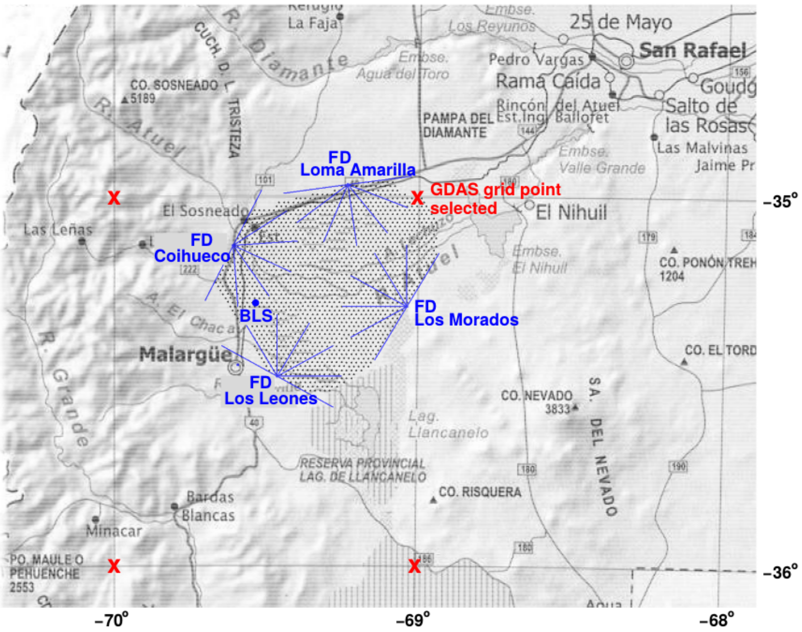

Figure 3.14: The GDAS grid points nearest to the Pierre Auger Observatory (red crosses). The most suitable grid point lies to the north-east of the array at (35°S,69°W). Image taken from [21].

are extracted from the archived GDAS data at the chosen grid point. Temperature, pressure, and relative humidity are readily available, while other relevant parameters, such as air density and atmospheric depth, must instead be calculated. In doing so, relative humidity is converted into water vapour pressure, which is a more fundamental quantity (the Magnus formula, see Equation 2.9 in Section 2.3). For the simulation and reconstruction of cosmic ray air showers, the atmospheric profiles are extended to 100 km using the 1976 U.S. Standard Atmosphere. The profiles are implemented into an Auger $\overline{\text{Off}}\underline{\text{line}}$ atmospheric database and the details are discussed in [21].

3.6.3 Comparing the GDAS data to local measurements

The applicability of data from the GDAS model for the Observatory has been explored in detail [21][75]. Measurements from both radio soundings and ground-based weather stations have been found to support the validity of the model [21][76]. In particular, the descriptions of the atmosphere at heights above ∼ 5 km a.s.l. have been shown to be in excellent agreement with observations from local radiosonde launches (see Figure 3.15). There are, however, some discrepancies between the GDAS data and the sonde measurements that become apparent near the ground. This is most noticeable in the temperatures predicted by the model, which have been found to consistently underestimate the observed temperatures (refer to Figure 3.15). As discussed in [21], a possible source for this disagreement may be that the model, which is designed to forecast atmospheric conditions on a global scale, does not properly describe very local environments—in particular, the influence that the neighbouring Andes mountains may have on the climate above the array. In addition, the model may not suitably describe the heated surface of the Pampa Amarilla plains. Despite these disagreements, further investigations have found the model to agree reasonably well with surface level measurements made by the several weather stations operating at Auger [21].

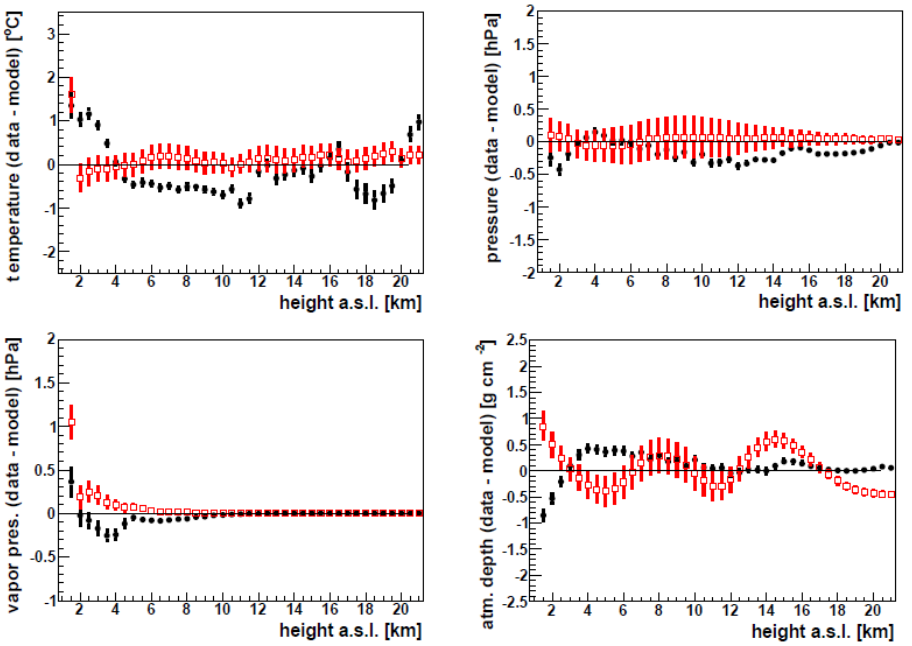

Figure 3.15: Differences between local radiosonde measurements and predictions made by the GDAS model (black dots) and monthly mean profiles (red squares) for several atmospheric quantities. Image taken from [21].

In a similar manner, I have also examined the accuracy of GDAS by comparing the 3-hourly predictions from the model to simultaneous weather station observations that provide a truth. As an example, Figure 3.16 shows GDAS minus Loma Amarilla weather station data for several atmospheric state variables. All available data from 2009 have been used. Figure 3.16a

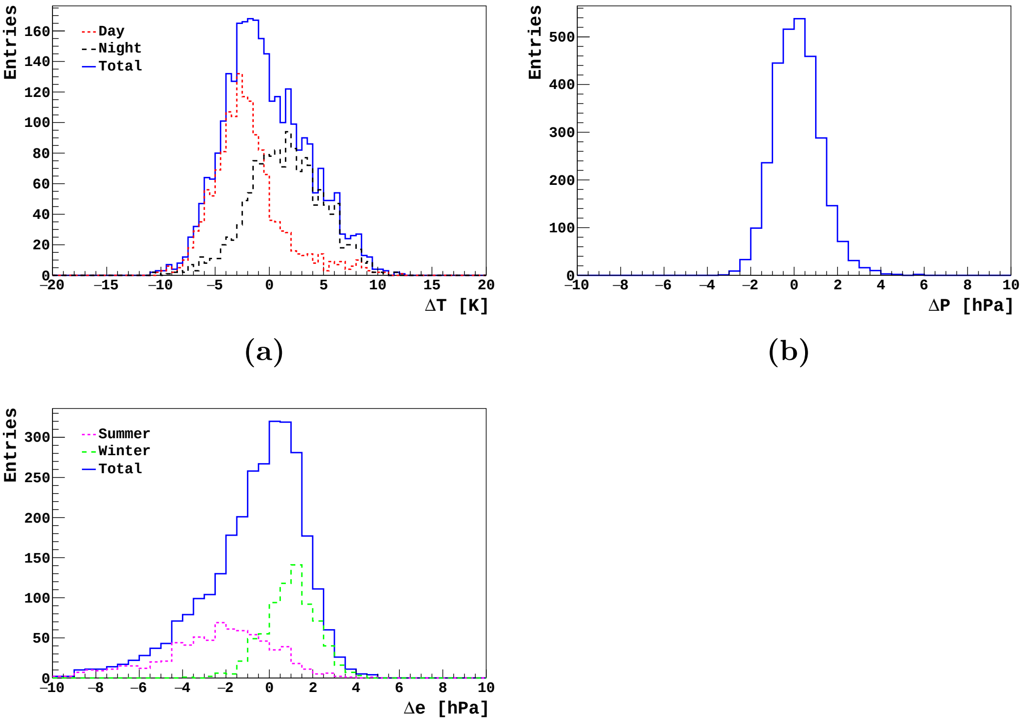

(a) **(b)**

(c)

Figure 3.16: Distributions of predictions from the GDAS model minus measurements made by the Loma Amarilla weather station. Data covers all of 2009 and shows differences in **(a)** temperature, **(b)** pressure, and **(c)** water vapour pressure. On average, the GDAS model underestimates ground temperatures during the day. Similarly, water vapour pressures are typically underestimated in Summer.

shows a distribution of temperature differences that has a mean of −0.3 K and a standard deviation of 3.8 K, indicating that GDAS describes surface temperatures reasonably well. The same can be said for air pressure (Figure 3.16b) and water vapour pressure (Figure 3.16c). However, by separating the temperature distribution into night (between 0 and 10 UTC) and day, the shortcomings of the model become more apparent. With a mean temperature difference of 1.5 K and a standard deviation of 3.5 K during the night, in comparison to a mean of −2.0 K and a standard deviation of 3.3 K throughout the day, it becomes clear that the model is prone to overestimating night-time surface temperatures. This behaviour can similarly be observed when comparing water vapour pressures predicted during Summer (Dec–Feb) and Winter (Jun–Aug) months. The results suggest that GDAS has a tendency to underestimate humidity in Summer. As noted in [21], calculating water vapour pressure depends strongly on air temperature (see Section 2.3) and could be more sensitive to the local effects surrounding the weather stations.

One possible source for the separate temperature distributions in Figure 3.16a may be due to the placement of the weather station. While some stations are located far from the FD buildings, others are mounted near or on top of facilities so that a standardised measurement cannot be guaranteed [21]. We offer another possible explanation, that the forecast, which assimilates real data into the model, is more accurate at certain times of the day. This is evident in Figure 3.17a, which shows the mean temperature differences in 2009 that occur throughout the day. There is a clear diurnal variation in the accuracy of the model. Our findings indicate that, on average, the best agreement occurs at 0 and 12 UTC.

Meteorological institutions around the world perform synoptic measurements of the upper atmosphere by performing routine radiosonde launches. The standard observing hours are at 0 and 12 UTC [77]. In doing so, this provides an instantaneous snapshot of the atmosphere, which is an important process for weather forecasting. We believe that improvements to the GDAS model at these synoptic hours likely reflects the data assimilation process. Although this behaviour is clear in the temperature data, there is no evidence of diurnal variations in air pressure (Figure 3.17b) or water vapour pressure (Figure 3.17c). It is important to note that although comparisons have only been made to the Loma Amarilla weather station, this behaviour is reproducible for the remaining four Auger weather stations. The extent of the effects are comparable at each site. Likewise, similar results can be achieved when investigating data from other years. Although this discrepancy is not a concern for cosmic ray analyses, it may lead to defining incorrect atmospheric profiles in MODTRAN (a radiative transfer program) such as in Section 5.4.

Although the above has highlighted some of the irregularities found in the GDAS model's description of surface values, they are not of considerable importance for most applications. In fact, the differences between the model and each weather station are comparable to those found between the individual stations across the array (after correcting for height differences). Overall, as demonstrated in [21] the model provides a good description of the atmosphere. The Pierre Auger Observatory has since transitioned to using data from GDAS to describe the molecular atmosphere.

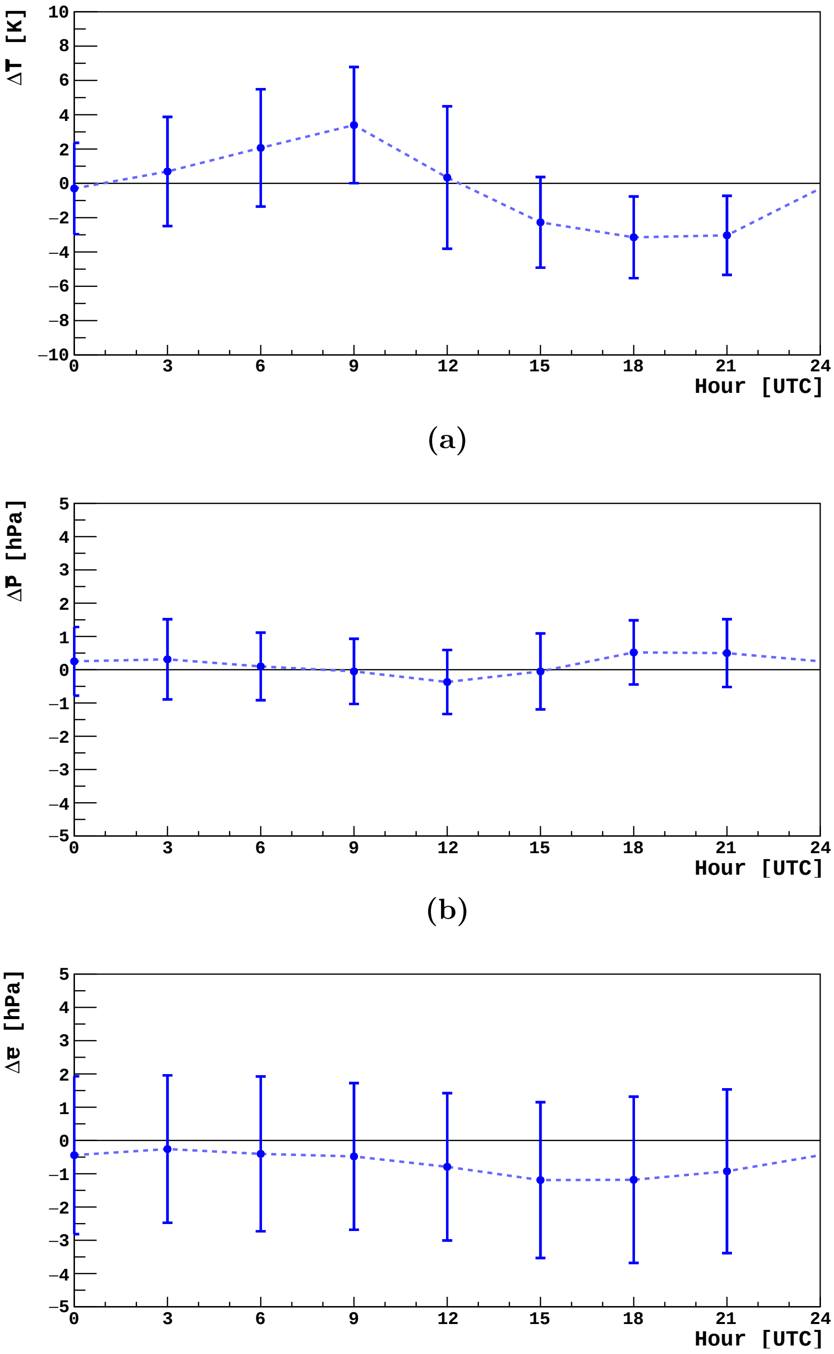

(a)

(b)

(c)

Figure 3.17: Data from the GDAS model minus measurements made by the Loma Amarilla weather station. The data covers all of 2009 and shows the mean differences for **(a)** temperature, **(b)** pressure, and **(c)** water vapour pressure throughout the day. Error bars denote the standard deviations of the differences.

3.6.4 Further applications

In our research we have benefited greatly from access to the height-dependent atmospheric data provided by the GDAS model. One application is the ability to determine the amount of moisture in the atmosphere at a given time (for the calculation, see Section 2.3.1). This is important to us, as both the infrared cloud cameras and single-pixel radiometers are sensitive to atmospheric moisture as well as clouds.

To illustrate and quantify the effects that cloud and water vapour can have on the infrared detectors, I have used the radiative transfer program MODTRAN. One of the features in MODTRAN is to allow the user to define an atmosphere by specifying the temperature, pressure, and humidity at different altitudes. This can be achieved with data obtained from GDAS and preliminary results are presented in Section 5.4.

In principle the GDAS atmospheric profiles can be used to determine the heights of cloud layers, in the same way as measurements from a radiosonde. This is investigated in Section 6.2 and results compared to lidar observations.

Field	**Units**	**Label**	**Data Order**
Pressure at surface	hPa	PRSS	S1
Pressure reduced to mean sea level	hPa	MSLP	S2
Accumulated precipitation (6 h accumulation)	m	TPP6	S3
u-component of momentum flux (3- or 6-h average)	$\mathrm{N\,m^{-2}}$	UMOF	S4
v-component of momentum flux (3- or 6-h average)	$\mathrm{N\,m^{-2}}$	VMOF	S5
Sensible heat net flux at surface (3- or 6-h average)	$\mathrm{W\,m^{-2}}$	SHTF	S6
Downward short wave radiation flux (3- or 6-h average)	$\mathrm{W\,m^{-2}}$	DSWF	S7
Relative Humidity at 2m AGL	%	RH2M	S8
U-component of wind at 10 m AGL	$\mathrm{m\,s^{-1}}$	U10M	S9
V-component of wind at 10 m AGL	$\mathrm{m\,s^{-1}}$	V10M	S10
Temperature at 2m AGL	K	TO2M	S11
Total cloud cover (3- or 6-h average)	%	TCLD	S12
Geopotential height	m	SHGT	S13
Convective available potential energy	$\mathrm{J\,kg^{-1}}$	CAPE	S14
Convective inhibition	$\mathrm{J\,kg^{-1}}$	CINH	S15
Standard lifted index	K	LISD	S16
Best 4-layer lifted index	K	LIB4	S17
Planetary boundary layer height	m	PBLH	S18
Temperature at surface	K	TMPS	S19
Accumulated convective precipitation (6 h accumulation)	m	CPP6	S20
Volumetric soil moisture content	fraction	SOLM	S21
Categorial snow (yes=1, no=0) (3- or 6-h average)	boolean	CSNO	S22
Categorial ice (yes=1, no=0) (3- or 6-h average)	boolean	CICE	S23
Categorial freezing rain (yes=1, no=0) (3- or 6-h average)	boolean	CFZR	S24
Categorial rain (yes=1, no=0) (3- or 6-h average)	boolean	CRAI	S25
Latent heat net flux at surface (3- or 6-h average)	$\mathrm{W\,m^{-2}}$	LHTF	S26
Low cloud cover (3- or 6-h average)	%	LCLD	S27
Middle cloud cover (3- or 6-h average)	%	MCLD	S28
High cloud cover (3- or 6-h average)	%	HCLD	S29
Geopotential height	m	HGTS	U1
Temperature	K	TEMP	U2
U-component of wind with respect to grid	$\mathrm{m\,s^{-1}}$	UWND	U3
V-component of wind with respect to grid	$\mathrm{m\,s^{-1}}$	VWND	U4
Pressure vertical velocity	$\mathrm{hPa\,s^{-1}}$	WWND	U5
Relative humidity	%	RELH	U6

Table 3.2: The meteorological fields contained in the GDAS data. The surface values 'S' are accompanied by a set of upper level values 'U' that are available at 23 different pressure levels. The highlighted fields have been implemented into an Auger Offline atmospheric database. Table adapted from [78].

Chapter 4

Camera Calibration Methods and Image Corrections

One of the University of Adelaide's contributions to the Pierre Auger Observatory has been to provide infrared cameras, which are used for night-time cloud detection. Our cameras are located on the roofs of each of the main fluorescence detector (FD) buildings (see map, Figure 3.1), and the current cloud cameras were installed progressively throughout 2013. The focus of the work featured in my dissertation has been on the most recent generation of cloud cameras operating at the Observatory (see Section 3.1.2).

Although the new cameras are a major improvement over the previous model (see Section 3.1.1), there are several techniques that can further enhance their image quality. In addition, because these new cameras are radiometric, the infrared measurements can be converted into temperature readings. Finding a way to convert a camera's signal into a temperature (i.e. a temperature calibration) has been one important aspect of my studies.

In this chapter, I present the methods used to calibrate our cloud camera systems, and some image processing techniques that can improve the image quality. One of the main obstacles, however, has been that the cameras had already been installed, and were collecting data at the Observatory. This meant that many of the routines had to be developed from our remote location, here in Adelaide. Another consideration has been that the techniques must be reproducible for each infrared camera.

4.1 Preliminary measurements

We have two infrared cameras at the University of Adelaide, which are the same model as those currently operating as cloud monitors at the Pierre

Auger Observatory (see Section 3.1.2). Both cameras have been fitted with an 18 mm lens that provides them with an approximate 30°×23° field of view. Our Adelaide thermal cameras not only act as potential replacements for the Auger cloud cameras, but they allow us to directly experiment with the instrumentation. For example, they have helped us to understand how the cameras respond to different software settings, as well as how the environment can affect their behaviour.

The experiments that I initially performed were to investigate the response of a camera to the temperatures of objects in the field of view. The major goal of these studies was to find a conversion from the digital output of a pixel (in ADC counts) into a temperature reading (i.e. a temperature calibration). In a laboratory environment, calibrations are typically performed using an experimental black body (an emissivity near 1) at a known temperature. This is because the radiometric properties of such an object are well-defined (see Section 2.1).

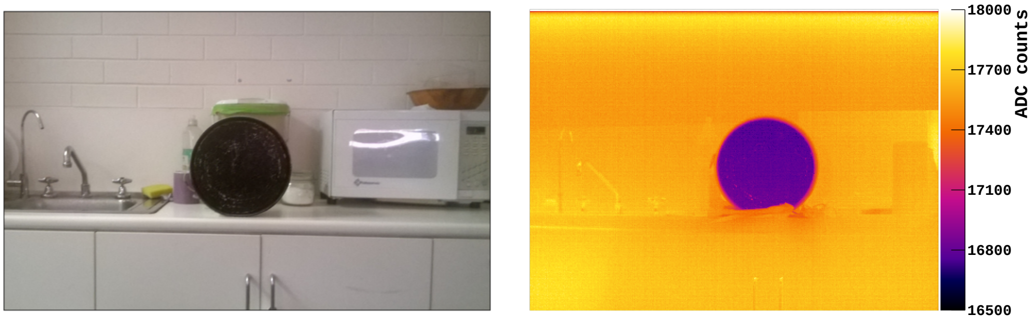

(a)

(b)

Figure 4.1: **(a)** A picture of the dark rounded aluminium plate (experimental black body) at visible wavelengths. **(b)** A false colour image of the same object at infrared wavelengths, as captured by one of our thermal imaging cameras. The temperature of the object was 280.8 K, and stands out against the warmer room background. The digital output of each pixel (in ADC counts) relates to the infrared brightness within the field of view.

The object that I used for these preliminary measurements was a dark rounded aluminium plate that resembles a black body. The aluminium has a thickness of 8 cm and a diameter of 23 cm, and has a rough surface coated in black matte paint, which helps to make it a better black body. In addition, a temperature sensor is embedded within the plate, which allows the absolute temperature of the object to be known. This particular aluminium plate had acted as a black body for previous studies at the University of Adelaide (for example [53]), and had also been used to calibrate our single-pixel infrared radiometers (discussed in Section 3.2). Figure 4.1a shows a picture of the object at visible wavelengths, whereas Figure 4.1b shows the same object at infrared wavelengths captured with one of our thermal imaging cameras.

4.1.1 The camera's response to scene temperatures

A simple way to test the response of the camera to different scene temperatures was to place the aluminium plate (experimental black body) within the camera's field of view. In order to observe the behaviour over a wide range of temperatures, the aluminium plate was cooled in a freezer prior to taking measurements. The temperature of the black body began at 260 K when data were initially recorded. Over the course of 2 hours, the object's temperature gradually warmed to about 280 K.

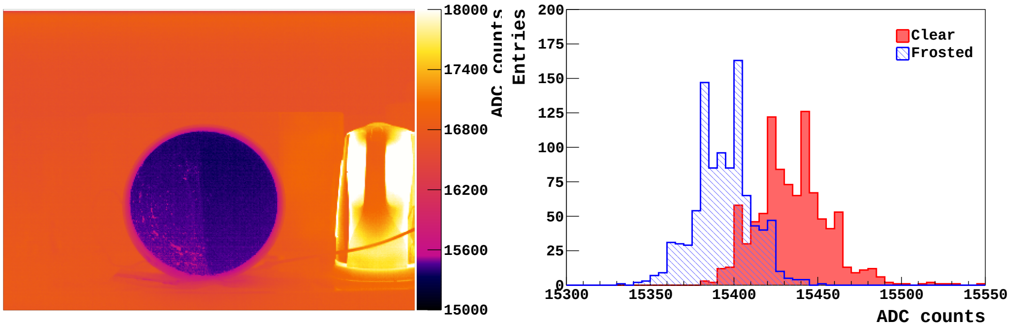

(a) (b)

Figure 4.2: (a) An infrared image of an aluminium plate, which resembles a black body. A layer of frost has accumulated on one side of the object (to the right of the image). The temperature of the body is 265 K, however, the frosted side appears cooler (a lower ADC count) due to the lower emissivity of ice. **(b)** A distribution of some pixel values corresponding to the aluminium plate for the same image. The frosted side (blue) has a mean value of 15393 ADC counts, which is slightly less than the clear side (red) that has a mean value of 15435 ADC counts. Data were collected on July 16th, 2014.

A layer of frost would occasionally coat the surface of the aluminium plate. Given that ice has a lower emissivity than our experimental black body [79], the formation of frost was undesirable because it could impact on my results. A classic demonstration of how differences in emissivity affect infrared measurements can be illustrated with a Leslie cube (for example [29]), where the polished sides (at a low emissivity) of the cube appear less bright than the rough sides (an approximate black body). Similarly, I was already familiar with the effects that frost could have on my work from another experiment using the same aluminium plate (summarised in Figure 4.2).

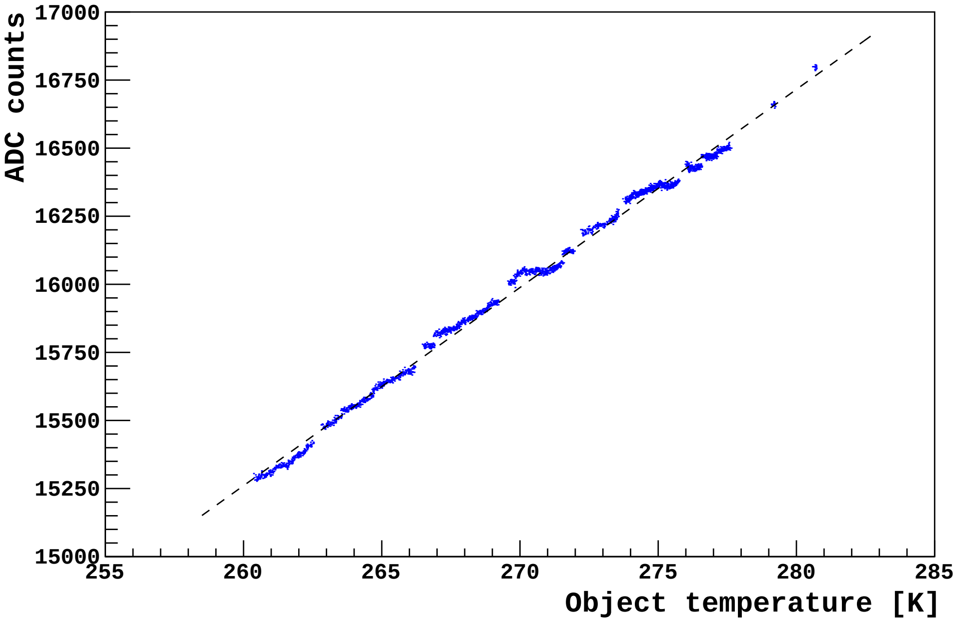

Figure 4.3: Measurements made by our thermal imaging camera for different scene temperatures, which were collected by placing a black body of a known temperature in the field of view. The body began at 260 K and gradually warmed to about 280 K. The camera's output (in ADC counts) relates to the body's infrared brightness, and governed by the object's temperature (Stefan-Boltzmann law). A linear relationship (dashed line) is a good approximation over this temperature range. The temperature of the camera's sensor can also affect the camera's response, but this remained fairly constant throughout this particular experiment. The sensor had a mean temperature of 324.6 K and a standard deviation of 0.5 K. Data are derived from 2590 images captured on June 19th, 2014.

Given the above, I removed frost from the aluminium object at various times throughout this particular analysis. The thermal images and aluminium temperatures were generally recorded at 2 second intervals, except towards the end of the experiment when the temperature of the aluminium varied slowly with time (Newton's law of cooling: the rate at which the temperature changed was approximately proportional to the temperature difference between the body and its surroundings). An average pixel value (in ADC counts) for the aluminium plate was determined in each image. I calculated the mean by averaging an 11×11 pixel array at the centre of the image of the object. The standard deviation of those pixel values were typically ~ 10 ADC counts. Subsequently, a calibration curve can be created by plotting the camera's output for the different scene temperatures. The results are presented in Figure 4.3.

It can be seen from the data in Figure 4.3 that there is a correlation between the camera's output and the temperature of the black body. A relationship can be expected since the infrared flux the camera receives depends on the temperature of the emitting object. Although the flux is proportional to the fourth power of the object's temperature (Stefan-Boltzmann law), these findings suggest that a linear relationship is appropriate over this small temperature range. Taken together, the camera's digital output (in ADC counts) of an object (at a temperature T_{obj} in Kelvin) can be converted into a temperature reading (i.e. a temperature calibration) using the following expression:

$$\mathrm{ADC} = m \times T_{obj} + C\,, \tag{4.1}$$

where $m = 72.8$ ADC counts per K, and $C = -3659$ ADC counts for this particular experiment, which was obtained by a linear fit to the data in Figure 4.3.

A simple linear relationship can also be applied to our Auger cloud cameras, given that they view a similar temperature range. However, the pixel output is influenced by other factors besides scene temperatures, for instance, the camera settings (such as the image span and image level, see Section 3.1.2.2), and the temperature of the camera's sensor. Although the settings can be managed using software, we cannot directly control the sensor temperature.

4.1.2 The camera's response to sensor temperatures

One quantity that has been found to affect our measurements is the temperature of the camera's sensor (i.e. the focal plane array, see Section 3.1.2), which we monitor, and will commonly be referred to as T_{cam} throughout my dissertation. The sensor temperature tracks with the ambient temperature, and is often found to be 25 K warmer than the local environment (due to an internal heater). This can be problematic for the Auger cloud cameras because they are located outside of the fluorescence detector buildings, where the temperatures are not controlled. On a typical night, T_{cam} changes by about 8 K. Annual variations in T_{cam} are even more substantial, ranging from about 290–320 K. Given that a camera's response can vary throughout a night (or year), understanding how the sensor temperatures can impact on our data is particularly important for a proper temperature calibration. Hence, some modification to Equation 4.1 is necessary to compensate for the effects.

In Section 4.1.1, the camera's response to different scene temperatures was investigated (for a constant T_{cam}). Whereas for this particular analysis, we instead want to vary T_{cam} for a fixed scene temperature. In order to cover a wide range of sensor temperatures (similar to what we observe with the Auger cameras), the camera (switched off) was cooled in a fridge prior to taking measurements. The temperature of the camera's sensor began at about 277 K when data were initially collected. Over the course of 75 minutes, T_{cam} progressively warmed to 320 K.

To see how our infrared measurements and temperature calibration (see Equation 4.1) are influenced by T_{cam}, two different scene temperatures were chosen for this study. One was the aluminium plate (experimental black body), which was stored in a freezer. The aluminium had an average temperature of 259.4 K which remained fairly constant (a standard deviation of 1.6 K) throughout this experiment. The second was a brick at room temperature, which was also placed within the camera's field of view. A brick was chosen to represent a black body at room temperature because it has a rather high emissivity (about 0.9 [80]). Using a hand-held infrared thermometer, the temperature of the brick was recorded 7 times during the experiment. The mean temperature of the brick was found to be 292.3 K and the measurements had a standard deviation of 0.4 K.

One of the challenges in my experiment was to maintain the temperature of the aluminium plate. Whenever I was ready to capture an image, I would briefly open the door to the freezer and expose the object. A total of 27 measurements were performed for this analysis. Average pixel values (in ADC counts) for the aluminium plate and the brick were determined for each measurement. I calculated the values by averaging a 21×21 pixel array at the centre of both objects. Subsequently, the camera's digital output for the two scene temperatures could be plotted against T_{cam}, and the results are presented in Figure 4.4.

What stands out in Figure 4.4 is the parabolic-like response of the camera, which can be seen for both scene temperatures. This is a rather interesting feature because it suggests there is a degeneracy in the camera's signal (i.e. for a given scene temperature, the output for a T_{cam} at 310 K will be the same at 320 K). Another implication is that the camera's response is significantly influenced by the sensor temperature, particularly when T_{cam} deviates from 315 K (the approximate minimum of both features in this figure). The ra-

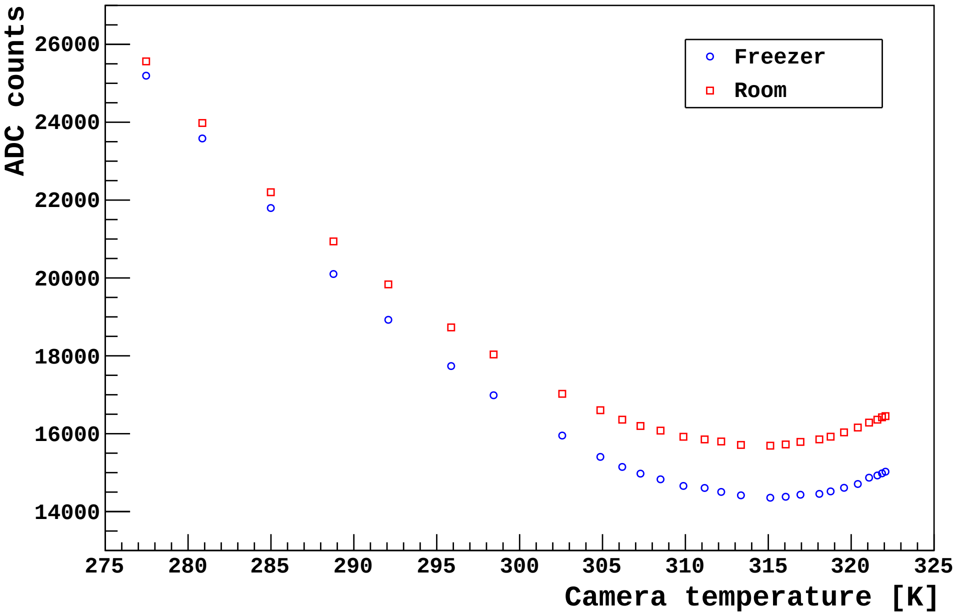

Figure 4.4: Measurements made by our thermal imaging camera at different camera sensor temperatures (T_{cam}). Two scene temperatures were viewed by the camera. One was a cold black body (an aluminium plate) at 259.4 ± 1.6 K, which was stored in a freezer. The second was a warm black body (a brick) at 292.3 ± 0.4 K, which was at room temperature. Data are derived from 27 images captured on July 16th, 2014.

pidity of the changes that are observed in the digital output emphasises why we wish to compensate for the temperature of the sensor. A similar response can also be found in other infrared camera sensors [49][81][82], and correcting for the behaviour has been the subject of previous studies (for example [83]).

Figure 4.4 demonstrates that T_{cam} can influence our measurements, however, it does not directly show how we might modify our temperature calibration (Equation 4.1). In this experiment, two scene temperatures with a difference of about ~ 30 K were chosen. We know from our previous experiment (see Section 4.1.1) that the camera's response to scene temperatures is approximately linear over this temperature range (see Figure 4.3). Hence, these two points are sufficient to perform a temperature calibration.

Figure 4.5a shows an example of a two-point temperature calibration when T_{cam} was at 322.1 K. One point is the black body (the aluminium plate) stored in the freezer, which was at 262.2 ± 0.1 K for this particular measurement. The second point is a black body (the brick) at room temperature, which was approximately $292.3 \pm 0.4\,\mathrm{K}$ for all 27 measurements. The camera's digital output was, on average, recorded to be 15024 ± 17 ADC counts for the first point, and 16451 ± 14 for the second point. It follows that for this particular observation, a two-point temperature calibration (given by Equation 4.1) would have a slope $m = 47.4 \pm 1.0$ ADC counts per K, and an offset $C = 2593 \pm 286$ ADC counts. This calibration curve is indicated by the dashed line in the figure. However, this was only for a single measurement. The same procedure can be extended to the remaining 26 observations. Subsequently, a relationship between the two-point calibrations and the temperatures of the sensor can be investigated.

Both Figure 4.5b and Figure 4.5c reveal that the two-point temperature calibrations change depending on the temperature of the camera's sensor. Figure 4.5b shows a steady increase in the calibration slopes at warmer sensor temperatures. In fact, a variable slope (which relates to a temperature difference) suggests that the temperature resolution of our camera changes. This behaviour may also be observed in Figure 4.4, given that the separation (in ADC counts) between the fixed scene temperatures changes with the camera's temperature (becoming narrower when T_{cam} is cooler). In addition, Figure 4.5c shows that the calibration offsets can vary significantly with the sensor temperature, and explains the parabolic-like features that were previously seen in Figure 4.4. Taken together, these findings suggest that both the slope $m(T_{cam})$ and offset $C(T_{cam})$ of a temperature calibration depend on the camera's temperature. Hence, it would be more appropriate to express Equation 4.1 by the following:

$$\mathrm{ADC} = m\left(T_{cam}\right) \times T_{obj} + C\left(T_{cam}\right) \,, \tag{4.2}$$

where the camera's digital output (in ADC counts) not only depends on the temperature of an object in the field of view (T_{obj} in Kelvin), but also the temperature of the camera's sensor (T_{cam} in Kelvin).

Because we already monitor the temperature of the camera's sensor, Equation 4.2 can be used to convert the camera's signal into a temperature reading provided that both a slope and an offset are known. A parameterisation for $m(T_{cam})$ and $C(T_{cam})$ will be addressed when I determine the temperature calibrations for the Auger cloud cameras in Section 4.5.

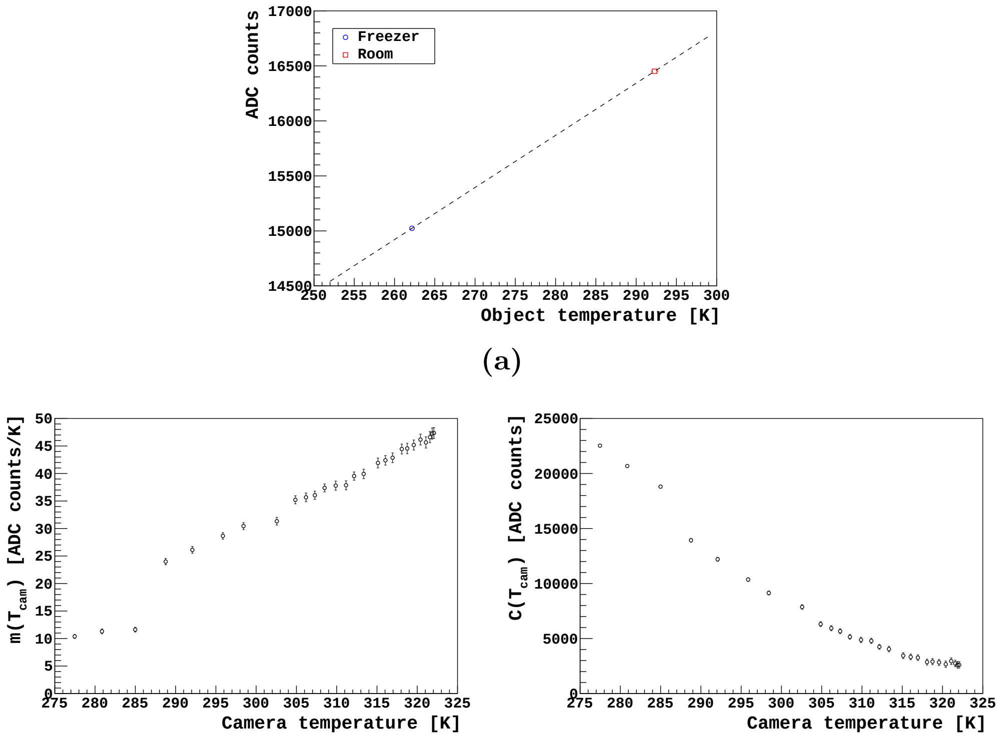

(a)

(b)

(c)

Figure 4.5: **(a)** An example of a two-point temperature calibration using a single image that was captured when the temperature of the camera's sensor (T_{cam}) was 322.1 K. One point is a black body at 262.2 ± 0.1 K (blue circle marker). The second point is a black body at 292.3 ± 0.4 K (red square marker). On average, the camera's digital output was 15024 ± 17 ADC counts for the first point, and 16451 ± 14 ADC counts for the second point. Consequently, this particular two-point calibration (dashed line) would have a slope $m = 47.4 \pm 1.0$ ADC counts per K, and an offset $C = 2593 \pm 286$ ADC counts. **(b)** Various two-point calibration slopes $m(T_{cam})$ calculated from 27 images captured at different sensor temperatures. **(c)** Various two-point calibration offsets $C(T_{cam})$ calculated from 27 images captured at different sensor temperatures. The inconsistent slopes and offsets for $T_{cam} < 285$ K are believed to result from condensation forming on the camera's lens. Errors are derived from uncertainties in the temperature and camera measurements.

4.2 A catalogue of clear night skies

In Section 4.1, preliminary temperature calibrations were produced for one of the infrared cameras we have access to at the University of Adelaide. The main goal of my research, however, has been to calibrate, and improve the image quality of the four cameras operating at the Pierre Auger Observatory. In a laboratory environment, these procedures are typically carried out using an approximate black body. However, the cameras in Argentina had already been installed, and were taking data at the Observatory. This meant that I had to develop my routines for calibrating the cameras, and correcting the images remotely. Moreover, the techniques had to be reproducible for each of the cloud cameras. The starting point for the majority of these routines was to select suitable images. These images were taken at times when the night sky was believed to be completely clear.

Clear nights have been selected for several reasons. Firstly, they can easily be distinguished from partially cloudy and overcast conditions. One particular use for clear skies has been in the development of temperature calibrations for our cloud cameras. These calibrations allow us to convert a camera's digital output (in ADC counts) into a sky temperature. Not only are temperatures a more meaningful unit for us to work with, but the conversion also compensates for the dependence of the digital output on the temperature of the camera's sensor (see Section 4.1.2). We calibrate the Auger cloud cameras by comparing clear sky images to temperature measurements made with the single-pixel radiometer installed at the Central Laser Facility (CLF), which has a similar spectral sensitivity. The method may be found in Section 4.5.

Another use for clear skies has been to create flat-field corrections for each camera. A flat-field is the response of a camera to a uniform source of illumination. Ideally, this would result in a flat image, where each pixel has the same intensity. However, distortions in the optical path due to the lens, and variations in pixel-to-pixel sensitivity produce image artefacts. In addition, our cameras view the sky through a protective window. Not only do the lens and window distort light, but they also emit thermal radiation. This increases the brightness of our infrared images. Flat-field images are required to compensate for these effects. Within the spectral range of our cameras, the infrared brightness of a clear sky remains fairly constant at small zenith angles. Hence, thermal radiation from a clear night sky can provide our cameras with a uniform source of illumination. Determining flat-field corrections for the Auger cameras using overhead clear sky images may be found in Section 4.3.

Although the infrared brightness (or effective temperature) of a clear sky varies slowly across small zenith angles, it increases rapidly when viewing angles near the horizon. This is a consequence of the rapidly increasing optical depths of water vapour and carbon dioxide, which are the main emitting gases within the band-pass of our cameras (see Section 2.2.2). Also at these wavelengths, clouds typically have effective temperatures that are warmer than a clear sky. An important step to identifying clouds in our thermal images is to understand how clear sky temperatures change with zenith angle. Parametrising these changes using data from our infrared cameras may be found in Section 5.1.

Given the above, it was useful to produce a catalogue of times when the night sky was believed to be clear. These times were selected by visually inspecting the full-sky mosaics taken by our cloud cameras. The mosaics are captured at 15 minute intervals, and provide a general description of the cloud cover at each fluorescence detector site. In sequence, the movement of clouds above the array can also be seen. Figure 4.6 shows an example of using information from the cameras to search for clear skies to record in my catalogue. Although the skies at the Los Leones and Los Morados sites are observed to be clear, partial cloud cover can be seen at Loma Amarilla and Coihueco. In my analyses, I am particularly interested in measurements when all operating cloud cameras agree on a clear sky. Consequently, this particular time is not recorded in my catalogue.

Using cloud camera data from June 2014 to February 2016, the catalogue I have constructed is presented in Table 4.1. In addition to times when the sky was believed to be clear, the number of cloud cameras operating N_{cam} was also recorded. Cloud camera data within these selected time periods have been used for the remaining analyses throughout this chapter.

Our cloud cameras may suddenly turn off during data taking, often due to a loss of power. When this happens, the temperature of a camera cools to that of the local environment. When the power is restored, the camera's sensor gradually warms to $\sim 25\,\mathrm{K}$ above the ambient temperature. This can cause a thermal shock for the cameras, and can result in incorrect measurements [29]. The affected data are easily recognised, as the images gradually appear brighter until a steady camera sensor temperature has been reached, typically after several minutes. After October 12th 2014, the cloud camera software was updated to include pictures of the internal shutter (which can be controlled manually) at the beginning and end of each scan. The temperature of the camera's sensor (T_{cam}) when both shutter images are captured is also recorded. If T_{cam} varies suddenly throughout a scan, the camera will likely be affected by a thermal shock, and the data are not always reliable. For this reason, I improve the quality of data in my clear night sky catalogue by selecting particular measurements. Only cloud camera data where T_{cam}, for a given scan (recorded over 5 minutes or 15 minutes depending on the type of scan), varies by less than 2 K are used in my research.

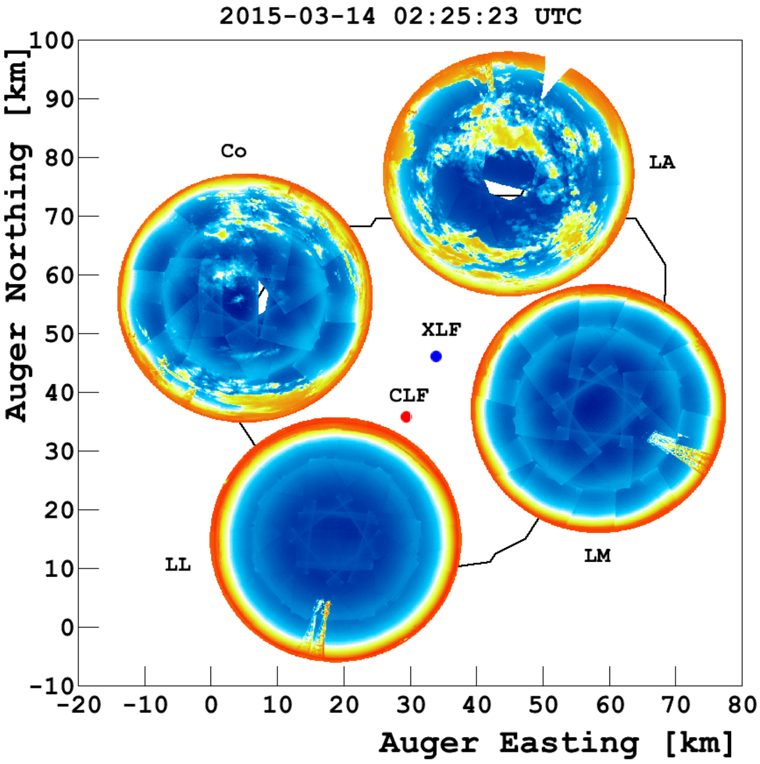

Figure 4.6: Contemporaneous full-sky mosaics from the four infrared cloud cameras operating at the Pierre Auger Observatory. The mosaics are captured within 4 minutes of one another, and are each comprised of 19 individual images. Clear skies can be seen at the Los Leones (LL) and Los Morados (LM) fluorescence detector sites, whereas partial cloud cover can be observed at Loma Amarilla (LA) and Coihueco (Co). The digital outputs are in ADC counts and relate to the infrared brightness in the cameras' fields of view. The output values are different for each camera. Each mosaic, however, ranges from low (blue) to high (red) infrared brightness. Clouds can be identified as appearing brighter than a clear sky. The mosaics are positioned on an outline of the Observatory, which also shows the locations of the Central Laser Facility (CLF) and eXtreme Laser Facility (XLF). With this arrangement, and in sequence, the movement of clouds across the array can be seen. This has been useful for producing a catalogue of times when the night sky was believed to be completely clear. Due to the partial cloud cover observed at two of the sites, this time is not recorded in the catalogue.

Table 4.1: A catalogue of nights at the Pierre Auger Observatory that are believed to be clear, based on the full-sky mosaics captured with our infrared cloud cameras. The table below shows the start and end times of these clear periods (in UTC). In addition, the duration of the period (in hours), and the number of cameras operating at the time N_{cam} are also shown. Camera measurements within these times have been used for the majority of the work found in this chapter.

Date	**Start**	**End**	**Duration [h]**	N_{cam}
2014-06-20	00:00:00	07:00:00	7	3
2014-06-27	00:00:00	10:00:00	10	4
2014-06-28	00:00:00	10:00:00	10	3
2014-06-29	00:00:00	10:00:00	10	4
2014-07-06	00:00:00	05:00:00	5	3
2014-07-07	01:00:00	09:00:00	8	3
2014-08-18	00:00:00	10:00:00	10	1
2014-08-19	00:00:00	10:00:00	10	1
2014-08-26	00:00:00	10:00:00	10	3
2014-11-14	01:00:00	08:00:00	7	3
2015-01-10	01:00:00	08:00:00	7	2
2015-01-12	01:00:00	08:00:00	7	2
2015-01-13	03:00:00	08:00:00	5	2
2015-01-21	03:00:00	08:00:00	5	2
2015-02-10	05:00:00	09:00:00	4	1
2015-02-11	01:00:00	09:00:00	8	2
2015-02-12	01:00:00	09:00:00	8	2
2015-02-13	03:00:00	09:00:00	6	2
2015-02-19	02:00:00	09:00:00	7	1
2015-02-24	03:00:00	09:00:00	6	2

Table 4.1 Continued:

2015-03-15	00:00:00	05:00:00	5	4
2015-03-16	00:00:00	09:00:00	9	3
2015-03-20	05:00:00	09:00:00	4	4
2015-03-29	00:00:00	05:00:00	5	3
2015-04-23	00:00:00	04:00:00	4	4
2015-04-26	00:00:00	10:00:00	10	4
2015-06-17	02:00:00	10:00:00	8	4
2015-06-26	00:00:00	03:00:00	3	4
2015-07-14	04:00:00	10:00:00	6	4
2015-07-23	00:00:00	05:00:00	5	4
2015-07-24	00:00:00	05:00:00	5	4
2015-08-18	00:00:00	05:00:00	5	4
2015-09-04	00:00:00	10:00:00	10	4
2015-09-05	00:00:00	10:00:00	10	4
2015-09-11	00:00:00	09:00:00	9	4
2015-09-12	05:00:00	09:00:00	4	4
2015-10-13	05:00:00	09:00:00	4	2
2015-11-14	00:00:00	08:00:00	8	4
2015-12-09	03:00:00	08:00:00	5	4
2015-12-11	03:00:00	08:00:00	5	3
2015-12-19	01:00:00	08:00:00	7	4
2015-12-20	01:00:00	06:00:00	5	4
2016-02-14	05:00:00	09:00:00	4	3

4.3 Flat-field corrections

A commonly used technique for improving the quality of digital images is to apply a flat-field correction (for example [84][85]). A flat-field is the response of a camera when it is exposed to a uniform source of illumination. This would ideally produce a flat image, where each pixel has the same intensity. In practice, however, the resultant images are typically non-uniform (see Figure 4.7). These artefacts can be corrected with various image processing methods. In this section, I outline the procedure we use to remove these undesired features in our infrared cloud camera images.

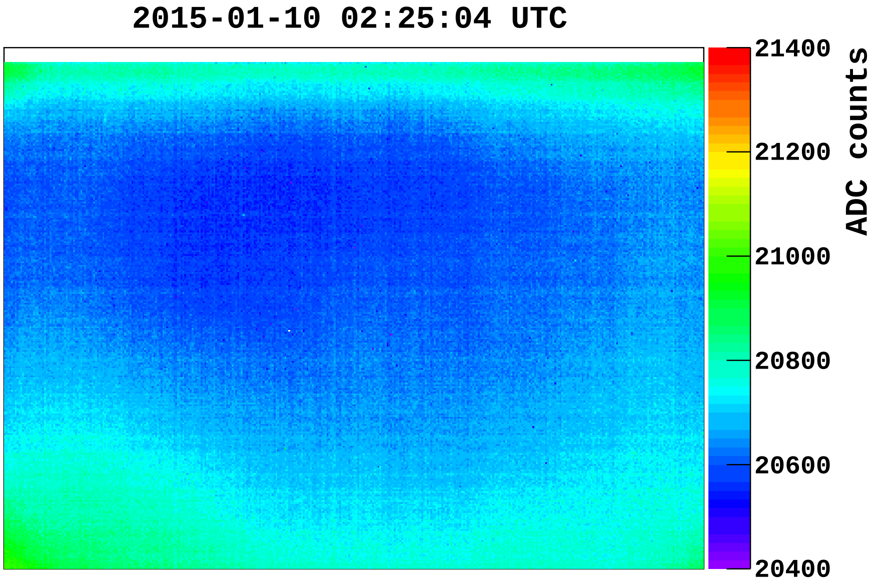

Figure 4.7: An example of a flat-field. The image has been created by exposing the Los Leones infrared cloud camera to a uniform source of illumination (the camera is pointed towards the zenith and measures thermal radiation from a clear night sky). This image should appear uniform, however, the optics of the system (the lens and window), variations in the pixel-to-pixel sensitivity, and other effects introduce image artefacts. Typically, our images appear warmer near the edges (a higher ADC count), which is where the optical path length through the emitting window material is greatest. The top four pixel rows are saturated and have been removed.

There are various reasons why a flat-field is typically non-uniform. A well-known example encountered in photography at visible wavelengths is vignetting, and is characterised as a reduction in brightness away from the image centre. There are several factors that contribute to vignetting. For instance, a natural reduction in brightness away from the optical axis [86], or the optical properties of the lens [87]. However, the situation is different with our infrared cameras, which are instead sensitive to thermal radiation. Not only does the lens distort incident light, but to the extent that light is absorbed, it also emits thermal radiation (Kirchhoff's law). As a result, our images typically appear warmer than expected near the edges. This is enhanced as our cameras also view the sky through a protective window, which emits thermal radiation as well. Hence, our cameras not only receive radiation emitted from an object within the field of view, but they also receive infrared flux contributions from both the lens and window. Consequently, our thermal images generally appear brightest near the edges, where the optical path length through the emitting window material is greatest.

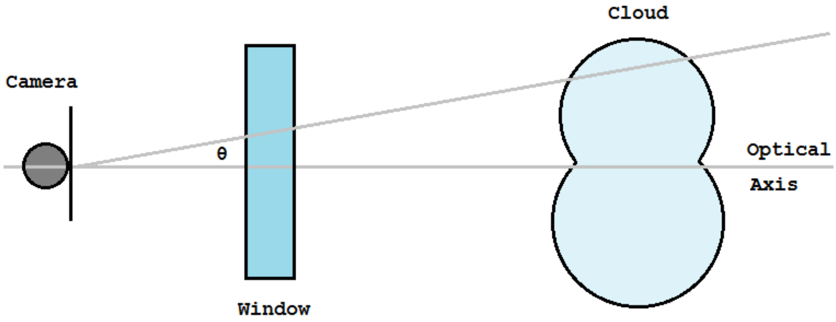

Figure 4.8: A system that shows a camera viewing a bank of cloud at an angle θ through a protective window. Both the cloud and window emit thermal radiation. The infrared flux received at the camera increases away from the optical axis, as the optical path length through the emitting window material is extended.

To demonstrate this concept, consider the system depicted in Figure 4.8. In this model, a camera views a cloud at an angle θ through a window. To simply the model, the camera lens is neglected, and only the cloud (at a temperature T_{cloud}) and window (at a temperature T_{window}) emit thermal radiation. First consider the radiant flux the camera receives along the optical axis (i.e. $\theta = 0°$ in the diagram). This can be expressed as:

$$\sigma T_{fov}^4 = \epsilon_{cloud}\sigma T_{cloud}^4 \tau_{window} + \epsilon_{window}\sigma T_{window}^4 , \tag{4.3}$$

where the effective temperature within the camera's field of view T_{fov}, cloud temperature T_{cloud}, and window temperature T_{window} are in Kelvin. Thermal radiation is emitted by the cloud with an emissivity ϵ_{cloud}, and the window with an emissivity ϵ_{window}. The fraction of the cloud's radiation passing through the window is given by τ_{window}, the transmissivity of the window. It follows that:

$$T^4_{fov} = \epsilon_{cloud} T^4_{cloud} \tau_{window} + (1 - \tau_{window}) T^4_{window} \,. \tag{4.4}$$

At an angle θ to the optical axis, the optical path length through the window is increased. Hence, a more general form of Equation 4.4 can be written as:

$$T^4_{fov}(\theta) = \epsilon_{cloud} T^4_{cloud} \tau_{window} \cos\theta + (1 - \tau_{window} \cos\theta) T^4_{window} \,. \tag{4.5}$$

To calculate the expected temperature increase at the edges of our camera images (i.e. where $\theta = 25°$, based on an approximate 50° field of view) when compared to the image centre ($\theta = 0°$), some values can be substituted into Equation 4.5. Estimating a value of $\tau_{window} = 0.8$ from the window transmittance at 11 µm (refer to Figure 3.3), and choosing other nominal values ($\epsilon_{cloud} = 1$, $T_{cloud} = 260\,\mathrm{K}$, $T_{window} = 300\,\mathrm{K}$), it can be shown that:

$$\frac{T_{fov}(25°)}{T_{fov}(0°)} \simeq 1.01 \tag{4.6}$$

It follows that when the field of view temperature at the image centre is 260 K, then the edges can be expected to appear ∼ 3 K warmer. This is approximately 300 ADC counts, and is similar to what we measure (see Figure 4.7). Equation 4.5 also shows that when the cloud and window temperatures are equal, then the field of view temperatures are the same regardless of angle to the optical axis. As the sky temperature (or cloud temperature) approaches the ground temperature (or window temperature) near the horizon, we can expect thermal emission from the window to have a lesser effect on our images. In addition, we can assume the infrared brightness due to the optics to be more pronounced in images of overcast skies (a warm effective temperature) compared to clear skies (a cool effective temperature). An example will be shown later in Section 4.3.2.

Although this model is basic, it demonstrates how the optics of our camera systems can influence our measurements. In reality, there are other factors that contribute to the non-flat thermal images we observe [88]. One factor is that not each pixel has the same response to light, and variations in the pixel-to-pixel sensitivity can lead to inhomogeneities across the camera's focal plane array. In addition, there is a thermal gradient across the array due to the placement of the camera sensor chip. Moreover, cool reflections of the camera off the lens and window can be captured in our images (the Narcissus effect) [89]. A straightforward approach to compensate for all these imperfections is to determine pixel corrections for the entire cloud camera system.

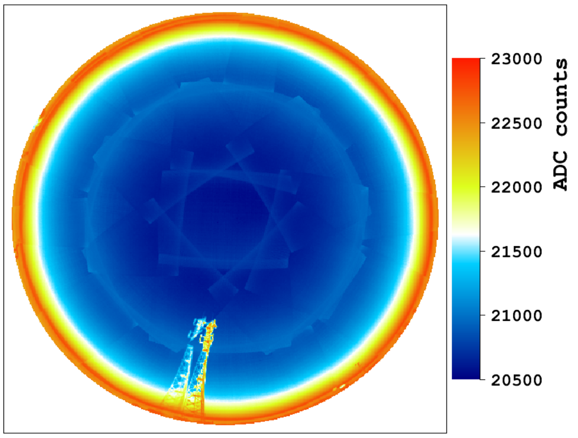

(a)

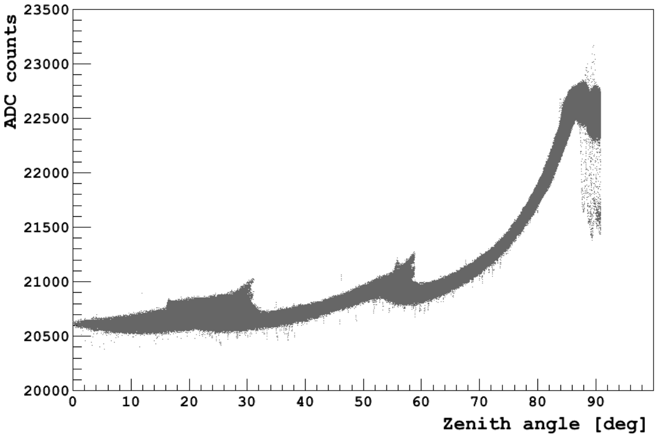

(b)

Figure 4.9: Measurements made by the Los Leones infrared cloud camera. The full-sky scan began at 2015-01-10 02:25:04 UTC. The scan consists of 19 images (10 are horizontal, 8 are elevated at 45°, and 1 is overhead) with some image overlap. **(a)** A full-sky mosaic that shows the images do not blend nicely together. This is mostly because each image appears brighter around the edges. Values in overlapping regions have been averaged. **(b)** The ADC counts for each pixel plotted against zenith angle. This should be smooth, but the brighter image edges creates scalloping features. Pixels corresponding to the radio tower have been removed for clarity. The same data with a flat-field correction applied will be shown later in Figure 4.13.

Although the need for a flat-field correction is not always obvious when looking at individual images, it becomes more apparent when connecting several images together (such as Figure 4.9). Figure 4.9a shows a full-sky mosaic that consists of 19 individual images, with some image overlap. The scan consists of 10 horizontal images, 8 images elevated at 45°, and 1 overhead image. One striking feature is that adjacent images do not blend together nicely. This is mostly because each image appears brighter around the edges. Besides affecting the full-sky mosaics, it can also be seen in Figure 4.9b, which shows the ADC counts for each pixel plotted against zenith angle. The infrared brightness (related to ADC counts) of a clear sky smoothly varies with zenith angle (see Section 5.1). However, rather harsh transitions are instead observed where the images overlap. This can be seen in the figure at 25° and 60° zenith angles. As we identify clouds that appear warm against a clear sky background, these warm image artefacts are problematic, and can lead to false cloud detections. This is the main reason we wish to remove these unwanted features.

4.3.1 Method

The unwanted features that have been discussed can be removed by looking at flat-field images captured with our cameras. The flat-field is the response of the camera to a uniform source of illumination. In the case of our infrared cameras, this source, for example, could be a black body that covers the entire field of view. Given the clear sky brightness within the spectral range of our cameras varies slowly with viewing angle, the camera output at zenith angles less than $\sim 30^\circ$ remains fairly constant. Hence, thermal radiation from a clear night sky represents another possible source of uniform illumination.

To determine the corrections for each pixel, an initial attempt was to use images of the internal shutter as a flat-field. The shutter acts as an approximate black body at a temperature similar to the camera's sensor, which is monitored. However, as the shutter is positioned between the sensor and the lens, the extent of the corrections due to the optics could not be determined with this method. To account for the lens and window, it was most useful to consider radiation emitted by an external source, such as the night sky.

Flat-field corrections have been calculated from a series of vertically captured images of clear night skies. The data are selected from a catalogue of skies I have constructed, which are listed in Table 4.1. Only data prior to May 5th, 2015 have been used. Each of the Auger cloud cameras has a unique flat-field correction. The corrections have been generated using 648 images from the Los Leones camera, 452 images from the Los Morados cam-

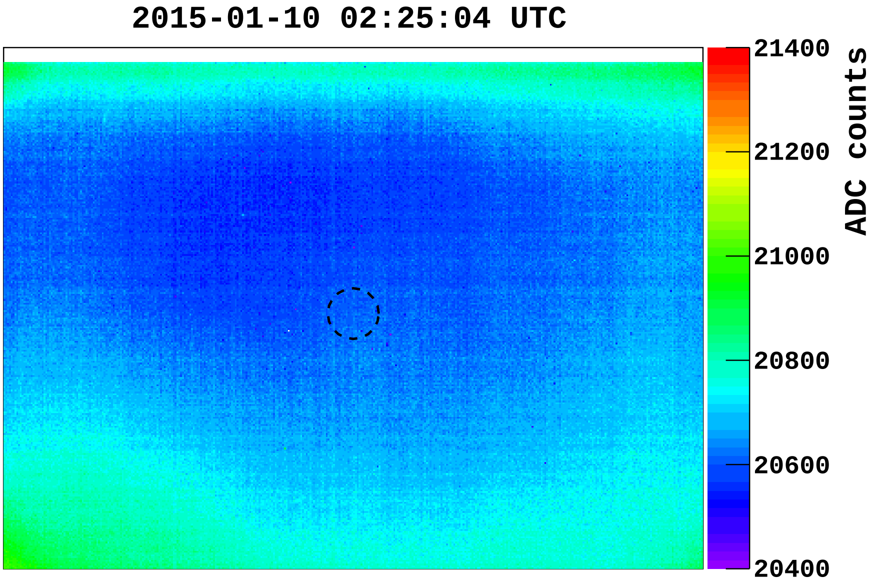

Figure 4.10: An example of an image used to determine a flat-field (the same as Figure 4.7) taken with the Los Leones infrared cloud camera. A dashed circle is drawn 2° around the image centre. The pixels within the circle have a mean of 20615 ADC counts, which serves as a reference (ADC_0). Each pixel will be adjusted to have the same brightness as the reference. The values of pixels within the circle have a standard deviation of 17 ADC counts.

era, 348 images from the Loma Amarilla camera, and 477 images from the Coihueco camera. To see how multiple images are used to generate a flat-field correction, first consider how a single image can be corrected.

Figure 4.10 shows an example of an image used to determine a flat-field (the same as Figure 4.7) taken with the infrared cloud camera at the Los Leones fluorescence detector site. In order to correct the intensity of each pixel (in ADC counts) a reference value is chosen (ADC_0). Each pixel will be adjusted to have the same brightness as the reference. Given that the entire image should be uniform, the intensity of any pixel can be chosen as a reference. It has been found that the image artefacts can be removed regardless of the choice of pixel. However, the choice does affect the overall image brightness. For this reason, a simple reference is at the image centre, where the warming effects due to the lens and window are expected to be minimised. To reduce the random error in ADC_0, it is more appropriate to average the brightness of multiple pixels instead of using a value from just a single pixel. I choose my reference to be the mean value of all pixels within 2° from the image centre ($\sim$ 600 pixels). These pixels are contained within

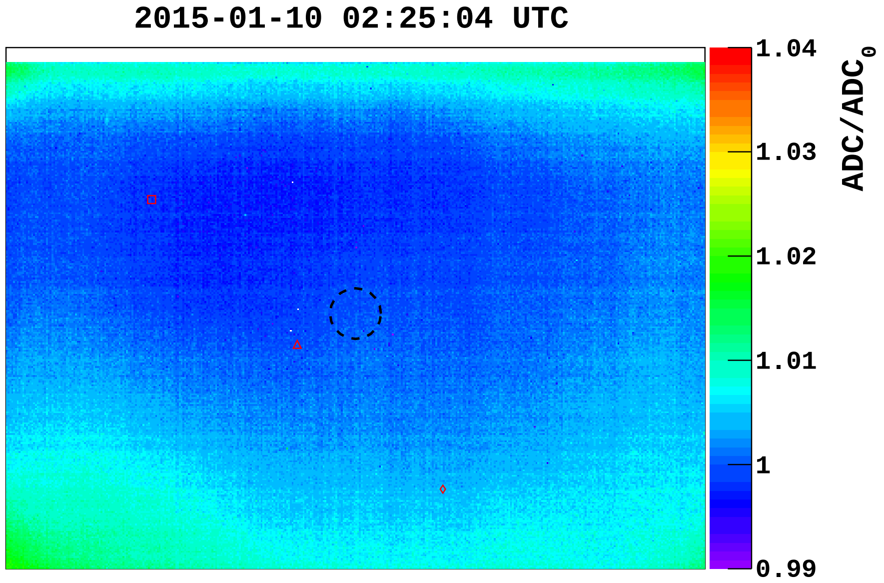

Figure 4.11: A relative comparison of each pixel to the reference (for the same image as Figure 4.10). The ratio of each pixel is calculated by dividing the value of each pixel (in ADC counts) by the reference (ADC$_0$). The pixels within the dashed circle have a mean of 1 by construction. Three red markers are also drawn. The square marker contains a pixel that has a (column,row) of (80,80) with a ratio of 0.997. Likewise, the triangle marker contains a pixel at (160,160) with a ratio of 1.000, and the diamond marker contains a pixel at (240,240) with a ratio 1.003. At this time the temperature of the camera's sensor (T_{cam}) was 320.58 K. The dependence of the ratios on the camera's sensor temperature (T_{cam}) for these three pixels are shown in Figure 4.12.

the dashed circle shown in Figure 4.10. In this example, ADC$_0$ is equal to 20615 ADC counts, and the pixel values have a standard deviation of 17 ADC counts.

After calculating the reference, a relative comparison of each pixel value to the reference (ADC/ADC$_0$) is evaluated. A relative comparison has been chosen as (anti)vignetting effects due to the lens and window are believed to be the main source of non-flatness in our images, and are characterised as a relative change in brightness away from the image centre. Through experimentation, however, developing the flat-field corrections in absolute terms (ADC − ADC$_0$) seemed to perform equally well. Figure 4.11 shows the ratio for each pixel. In addition, the pixels within 2° of the image centre are contained within the dashed circle, and have an average ratio of 1 by construction. It follows that if the ratio R of each pixel is known, a correction can be applied. That is:

$$\mathrm{ADC}_{ff} = \mathrm{ADC}_0 = \frac{\mathrm{ADC}}{R}\,, \tag{4.7}$$

where ADC and ADC_{ff} (or ADC_0) are the values of a pixel before and after the flat-field correction, respectively. In this example, each pixel will be adjusted to have a value of 20615 ADC counts (i.e. the reference). Although this calculation is circular when looking at how to correct a single image, the main objective is to determine ratios that can be applied to any image.

One quantity that has been found to affect the ratios is the temperature of the camera's sensor (T_{cam}). This is not surprising considering both the pixel output and temperature resolution depend on T_{cam} (see Section 4.1.2). In addition, there is an expectation that the infrared flux received by a camera in one of our systems depends on the temperatures of both the lens and the window (refer to Equation 4.5). Given that the lens and window temperatures track with the ambient temperature, and our camera sensors are typically 25 K warmer than the local environment, it follows that T_{cam} can also be used as a proxy for the lens and window temperatures.

To demonstrate there is a relationship between the ratios and T_{cam}, consider three of the pixels within the 384×288 image array. One of the pixels in Figure 4.11 is enclosed in the red square marker, and has a (column,row) of (80,80) and a ratio of 0.997. Similarly, the red triangle marker contains a pixel at (160,160) with a ratio of 1.000, and the red diamond marker contains a pixel at (240,240) with a ratio 1.003. The image was captured when the camera sensor temperature was 320.58 K. One way to test the dependence of these ratios on T_{cam} is to apply the same procedure to other flat-field images taken from my clear night sky catalogue. As the times that are recorded in the catalogue typically cover most of a night, and the dates are at various times throughout a year, the camera sensor temperatures span a wide range of typical operating values. The results are presented in Figure 4.12, and shows the ratios plotted against the temperature of the camera's sensor for the three chosen pixels.

For each pixel, I parameterise the relationship between the ratio R and the temperature of the camera's sensor T_{cam} by a linear function. That is:

$$R = p_1 \times T_{cam} + p_0\,. \tag{4.8}$$

This procedure can be extended to the remaining pixels that form the image array (384×288 total pixels). It follows that Equation 4.7 and Equation 4.8 can be used to apply a flat-field correction to each pixel (i.e. each pixel will have a value of ADC_{ff} after the correction). The extent of the correction

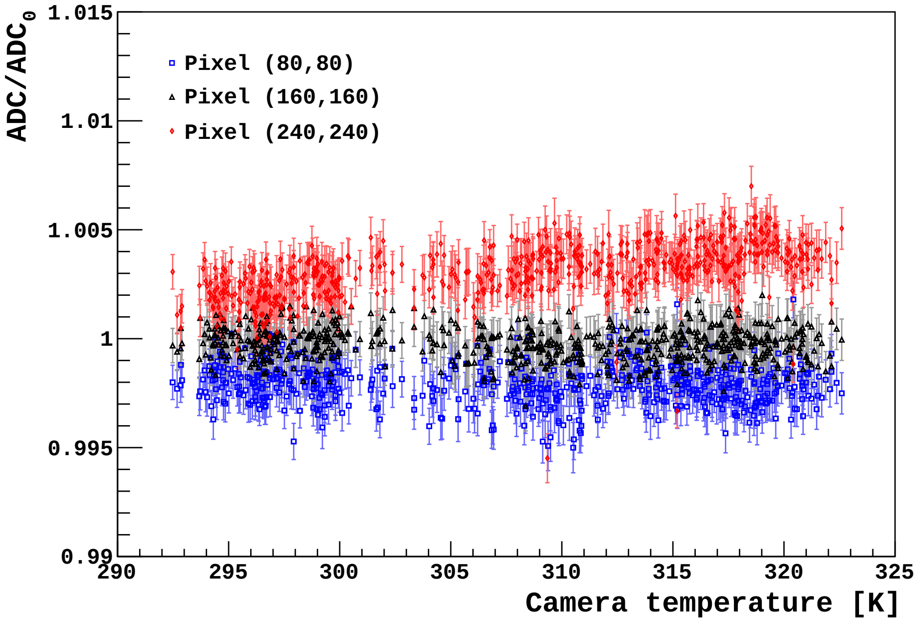

Figure 4.12: The ratio (ADC/ADC$_0$, or R in Equation 4.7) for three of the Los Leones cloud camera pixels (contained within the red square, red triangle, and red diamond in Figure 4.11) plotted against the temperature of the camera's sensor (T_{cam}). This shows that the extent of the flat-field corrections depends on T_{cam}. For each pixel, the dependence can be parameterised using a linear function (see Equation 4.8). The relationship has been determined from 648 individual flat-field images using a clear night sky. Error bars derived from the standard deviation of ADC counts in the pixels used to calculate the reference (ADC$_0$).

has a dependence on T_{cam}. As we already monitor the temperature of the camera's sensor, the remaining information needed to correct an image are the fit parameters p_1 and p_0. To this end, the parameters for each pixel are saved into a text file. This forms a template that can be read into a program when a flat-field correction is required. An example of using the template created for the Los Leones camera to apply a flat-field correction to images is shown in Figure 4.13. When compared to images without the correction applied (refer to Figure 4.9), it can be seen that the majority of image artefacts have been successfully removed. What stands out in particular is the smoother relationship between pixel ADC counts and zenith angle (see Figure 4.13b). This is especially useful for cloud detection, as it allows us to set a lower temperature threshold for selecting cloud pixels. This improves our sensitivity.

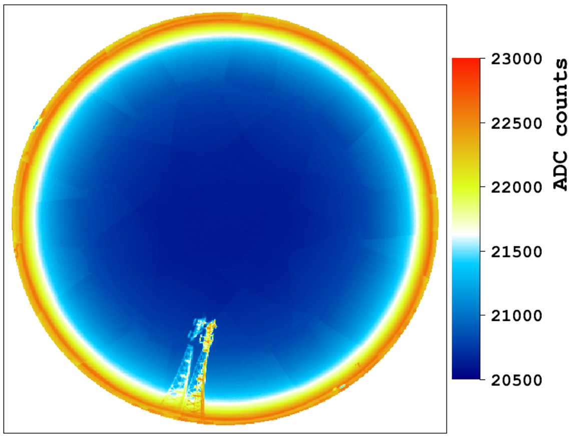

(a)

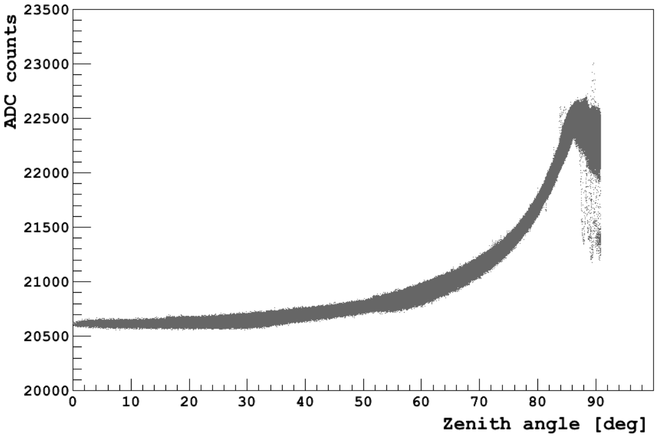

(b)

Figure 4.13: Measurements made by the Los Leones infrared cloud camera at 2015-01-10 02:25:04 UTC. A flat-field correction has been applied to remove several image artefacts. The full-sky scan consists of 19 images (10 are horizontal, 8 are elevated at 45°, and 1 is overhead) with some image overlap. **(a)** A full-sky mosaic. Values in overlapping regions have been averaged. **(b)** The ADC counts for each pixel plotted against zenith angle. Pixels corresponding to the radio tower have been removed for clarity. The same data without the flat-field correction have been shown in Figure 4.9.

This method of generating flat-field corrections for the Los Leones cloud camera can be successfully applied to the three remaining Auger cameras, which are installed at the Los Morados, Loma Amarilla, and Coihueco fluorescence detector sites. One interesting study then, is to compare the corrections we have developed for each camera. Given that the corrections depend on the camera sensor temperature, the comparisons have been evaluated for a T_{cam} of 300 K. The results are presented in Figure 4.14.

An interesting outcome of this study has been the differences in the flat-field for each of the cameras, and the magnitudes of the corrections that are required. While the shapes of the flat-fields for the Los Leones and Loma Amarilla cameras appear to be similar, those for the Coihueco and Los Morados cameras are revealed to be different. In particular, the radial increase in brightness away from the image centres is more pronounced. As a result, images captured with the Los Morados and Coihueco cameras also require a greater flat-field correction. In passing, we were already familiar with this when looking at data recorded with the Coihueco cloud camera, and it was one of the reasons we wished to develop a flat-field correction. Despite the apparent differences between cameras, another interesting study has been to determine the effectiveness of our flat-field corrections.

In order to assess the performance of each camera's flat-field correction, it is useful to determine the error (in ADC counts) in a pixel's value after the correction has been applied. One way of calculating this error is to compare the true value of a pixel (ADC_{true}), which has not been corrected, to a prediction that has been made using the template (ADC_{pred}). In each flat-field image, a reference value (ADC_0) can be calculated by averaging all the pixels within 2° of the image centre. Then a pixel's value can be estimated using Equation 4.7 and Equation 4.8. That is:

$$\mathrm{ADC}_{pred} = R \times \mathrm{ADC}_0 \,. \tag{4.9}$$

This can be applied to the N flat-field images (i.e. $N = 648$ images for the Los Leones camera) in the clear night sky catalogue. The root-mean-square error (RMSE) can be calculated for each pixel. This is given by:

$$\mathrm{RMSE} = \sqrt{\frac{\sum_{i=1}^{N} \left(\mathrm{ADC}_{pred} - \mathrm{ADC}_{true}\right)^2}{N}} \,. \tag{4.10}$$

The RMSE of each pixel for each of the Auger cameras can be found in Figure 4.15. A small error corresponds to a more effective flat-field correction. It can be seen that the errors after applying a flat-field correction are typically 20 ADC counts, which is approximately 0.2 K. Given that the digital outputs of our cameras are typically $\sim$ 20,000 ADC counts, the relative errors

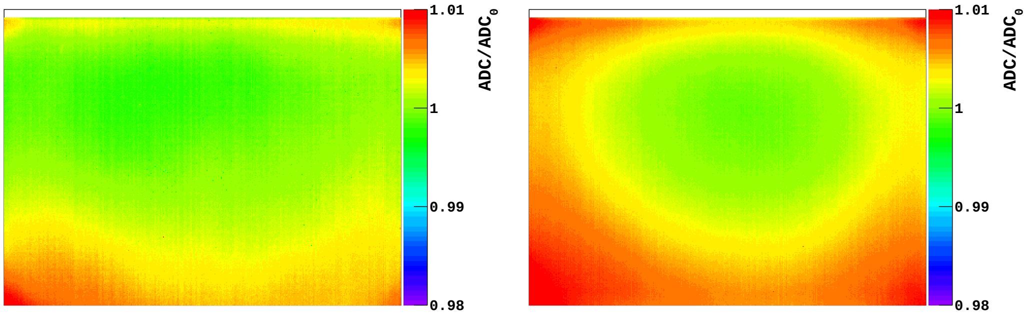

(a) Los Leones **(b)** Los Morados

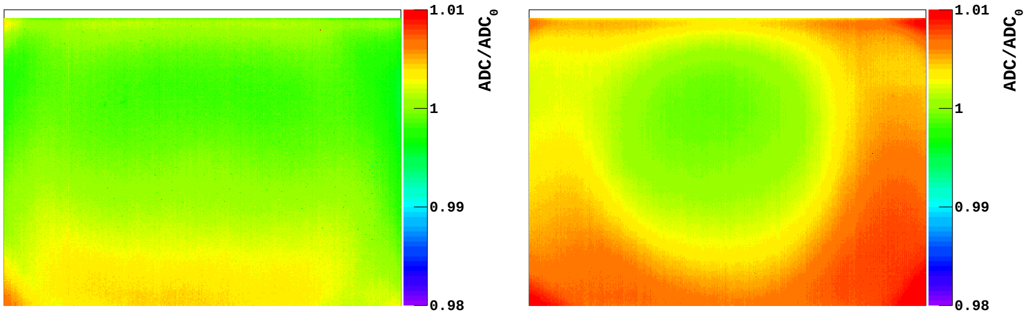

(c) Loma Amarilla **(d)** Coihueco

Figure 4.14: The flat-field correction templates for each of the Auger cloud cameras. The relative increase in brightness (ADC/ADC$_0$) of each pixel has been compared to a reference (ADC$_0$). The reference is the average value of the pixels within 2° of the image centre. Multiple clear sky images have been used to generate these templates. The extent of the corrections depends on camera sensor temperature (T_{cam}). To compare between the templates for each camera, the flat-field corrections have been adjusted to what they would be for a T_{cam} at 300 K.

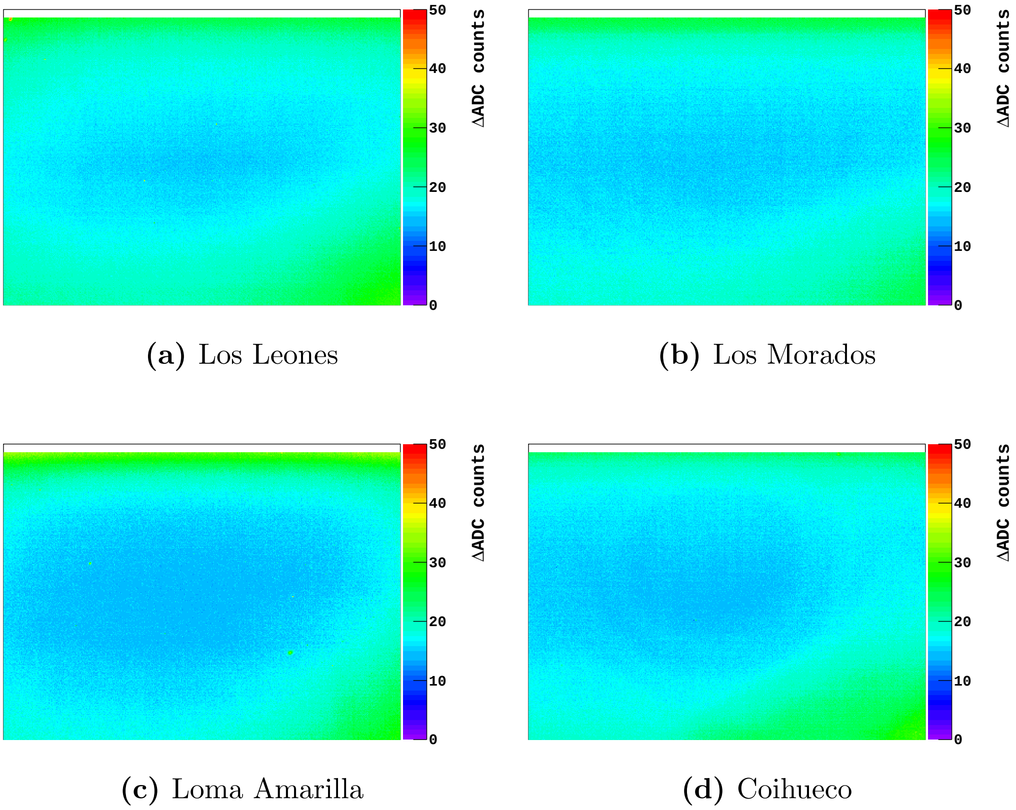

(a) Los Leones **(b)** Los Morados

(c) Loma Amarilla **(d)** Coihueco

Figure 4.15: The root-mean-square error (RMSE) of the pixels after a flat-field correction has been applied to the Auger cloud cameras. The errors (in ADC counts) have been calculated by comparing the value of a pixel to a prediction made using the flat-field correction (see Equation 4.9 and Equation 4.10). This is repeated for a set of flat-field images. The overall error after applying the correction is ~ 20 ADC counts, which is approximately 0.2 K.

associated with the flat-field corrections are suitably small for our purposes. Additionally, Figure 4.16 shows a histogram of ADC_{pred} minus ADC_{true} for one of the pixels. The distribution is approximately centred about the origin (a mean of 1 ADC count) and has a standard deviation of 18 ADC counts. Although this analysis was only for a single pixel, similar results can be obtained for the remaining pixels, and for our other cameras.

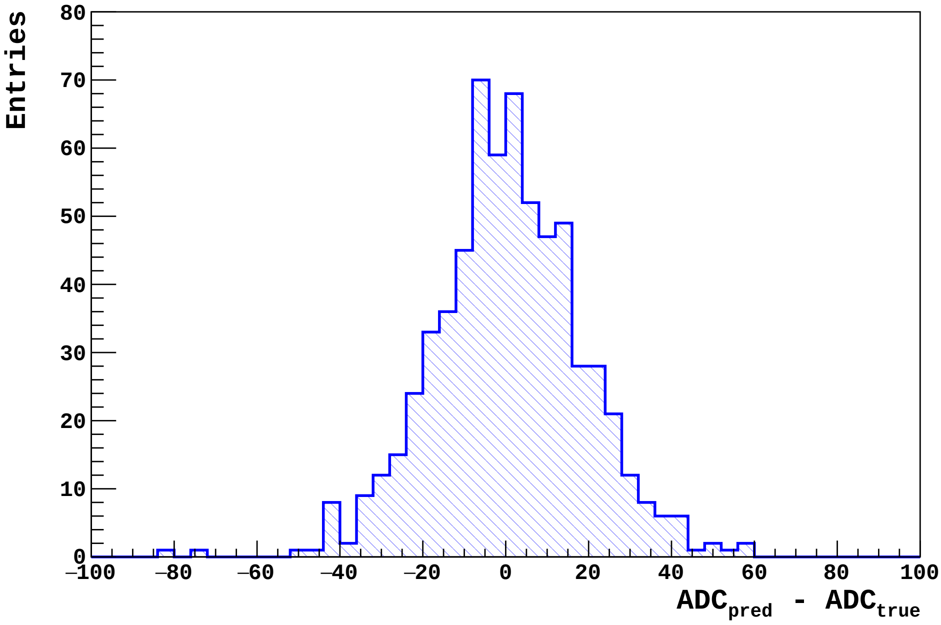

Figure 4.16: A distribution of predicted (ADC_{pred}) minus true (ADC_{true}) values for a single pixel using 648 images from the Los Leones camera. Predictions were made using the flat-field correction template and Equation 4.9. The true values were obtained from the original images without the flat-field correction. The distribution is approximately centred about the origin (a mean of 1 ADC count) and has a standard deviation of 18 ADC counts. This particular histogram was for the pixel at a (column,row) of (80,80), and similar results can be found for the other pixels (and other cameras).

4.3.2 Concluding remarks

In this section, flat-field corrections were developed for the four Auger cloud cameras. The main reason we wished to develop the corrections has been to improve our sensitivity to clouds. This is achieved by removing various image artefacts, which are mostly caused by the optics of our cloud camera units.

Our image processing technique uses thermal radiation from the zenith for a clear night sky. This can provide a camera with an approximately uniform source of illumination, and allows us to compensate for the non-uniformities that are observed in our images. Each pixel has a unique correction, which is determined by comparing the value of that pixel (in ADC counts) to some reference (ADC_0). Relative comparisons have been chosen (i.e. $\mathrm{ADC}/\mathrm{ADC}_0$) because (anti)vignetting effects due to the lens and window are believed to be the main reason why our images are not uniform, and are characterised by a relative change in brightness away from the image centre. However, as mentioned previously, absolute comparisons (i.e. $\mathrm{ADC} - \mathrm{ADC}_0$) seemed to perform equally well. In either case, the extent of the corrections depends on the temperature of the camera's sensor (see Figure 4.12).

It also seems reasonable for the flat-field corrections to depend on the temperature of an object within a camera's field of view. The digital output of a pixel not only relates to the sensor temperature, but also the scene temperature (refer to Equation 4.2). Moreover, our simple model also demonstrated that the infrared brightness of our images depends on the scene temperature (see Equation 4.5). Despite this, the small amount of scatter in Figure 4.12 suggests that scene temperatures do not have a major impact on the corrections for this set of data.

Through experimentation, the flat-field corrections had also been developed using temperatures rather than ADC counts. The effective temperature viewed by each pixel was calculated using our temperature calibration (see Section 4.5). In the same way, the temperatures (T in Kelvin) were compared to a reference (T_0) taken at the image centre. Both relative comparisons (T/T_0) and absolute comparisons ($T - T_0$) were investigated, and appeared to perform equally well. Interestingly, when temperatures are instead used, the flat-field corrections no longer depend on the camera sensor temperature (T_{cam}). This is because our temperature calibrations compensate for variations in the sensor temperature. In fact, the lack of a correlation with T_{cam} suggests our temperature calibrations are sensible. Although the flat-field corrections can be developed using temperatures, it is more straightforward to keep the pixel values in ADC counts.

Variations in scene temperatures do not appear to have a major impact on flat-field corrections using our clear night sky data. One way to investigate whether the temperature of field of view objects can affect the flat-field corrections is to consider much warmer sky temperatures. Warmer sky (or scene) temperatures can be achieved with overcast skies rather than clear skies. Thermal radiation emitted by completely overcast skies can also provide us with a uniform source of illumination, provided the cloud layer has no structure.

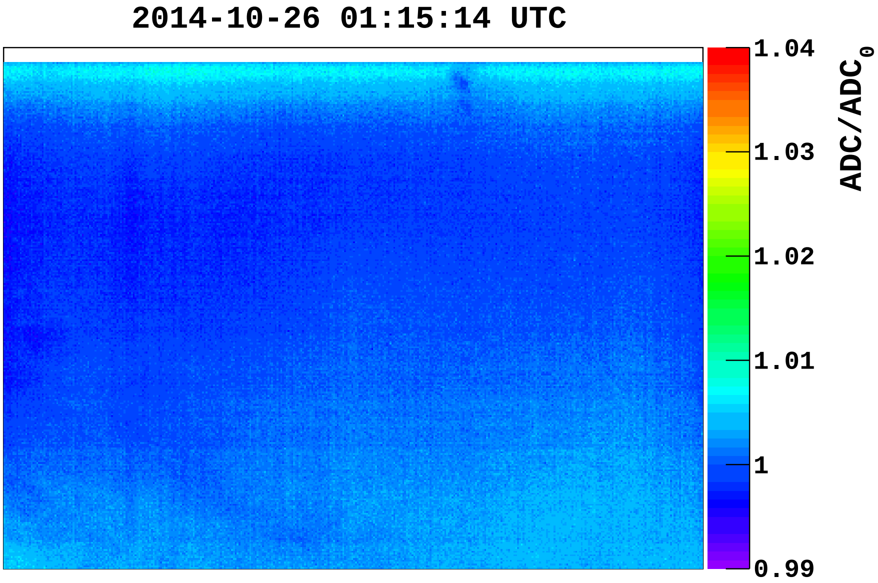

Figure 4.17: An example of a different flat-field taken with the Los Leones infrared cloud camera. Thermal radiation from an overcast sky (rather than the clear sky in Figure 4.11) instead provides the camera with a uniform source of illumination. The effective temperature of the overcast sky is $\sim$ 288 K (compared to $\sim$ 268 K for the clear sky). A relative comparison of each pixel value has been made to a reference. The reference is the average value of pixels within 2° of the image centre. At this time the temperature of the camera's sensor (T_{cam}) was 320.20 K. This is similar to the previous clear sky example, and so a direct comparison can be made. The pixel values for this warm (overcast sky) image require a smaller flat-field correction than the cooler (clear sky) image.

The response of the Los Leones cloud camera to an overcast sky may be found in Figure 4.17. Following the same procedure as before, a relative comparison of each pixel value has been made to a reference. The reference is the average value of all pixels within 2° from the image centre. In this example, the temperature of the camera's sensor is 320.20 K, which is similar to our previous example for a clear sky. Consequently, a direct comparison between both flat-field images can be made. The overcast sky has an effective temperature of ~ 288 K, whereas the clear sky has an effective temperature of ~ 268 K. One obvious difference is that pixel values in the warmer image require a smaller flat-field correction. Hence, applying a flat-field correction for a clear sky to an overcast sky can lead to an overcorrection (~ 100 ADC counts, approximately 1 K) in the images.

It was challenging to incorporate a dependence on scene temperature into the flat-field corrections. The reason for this difficulty is due to the rather harsh differences between clear and overcast flat-field images. A smooth transition for the corrections in terms of scene temperature could not be found. However, in practice, the flat-field corrections developed using only clear night skies perform sufficiently well, and remove the majority of image artefacts. Although there can be an overcorrection in overcast images, this does not impact on cloud detection.

In the past, I encountered a problem with the flat-field correction for the Coihueco cloud camera. Instead of removing image artefacts, the correction would introduce an unwanted feature by brightening the centre of the images. The problem was identified in November 2015, and would occur at times when the local humidity was high. We suspect the problem may have been caused by a build-up of moisture, which had been observed behind the protective window material. In order to remove the trapped water vapour, a desiccant sachet was placed within the cloud camera housing in December 2016. Since the placement of the desiccant, the issue seems to have been resolved. The performance of the flat-field correction for the Coihueco camera can be monitored in the future, particularly during the Summer months when more atmospheric water vapour is present.

4.4 Additional image corrections

In the previous section, flat-field corrections for the Auger cloud cameras were developed. The goal of the corrections was to improve the quality of our digital images by removing various artefacts. One major source that creates these artefacts is the optics of the system (the lens and window), which emits thermal radiation. This had been the main focus of Section 4.3. Other unwanted features appear in our images because of variations in the pixel-to-pixel sensitivity. Each pixel behaves as an individual infrared detector (a bolometer), and each detector has a separate gain and signal offset [29]. Consequently, a spread in detector responses also leads to non-uniformities across the camera's image sensor [90]. This typically results in fixed pattern noise [91], which degrades the image quality (for example [92]).

I have designed our flat-field corrections on a per-pixel basis, where each pixel has a unique correction. The procedure compensates for the majority of non-uniformities observed in our images. This is an image processing technique that is applied after measurements have been taken. The infrared cameras, however, also automatically perform their own corrections by recalibrating. This periodically occurs when the cameras are operating.

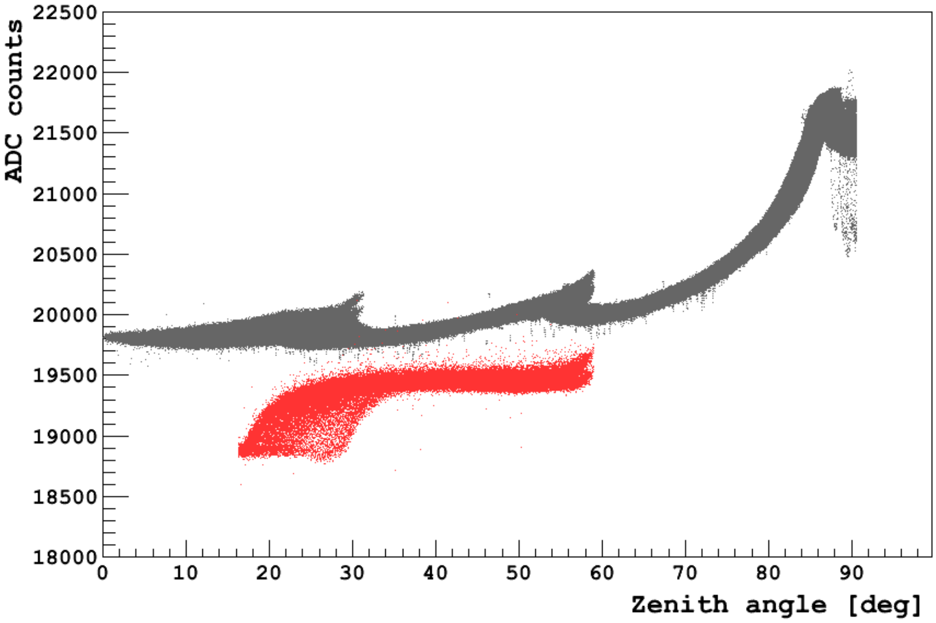

Figure 4.18: Measurements made by the Los Leones infrared cloud camera before applying a flat-field correction to the images. The full-sky scan began at 2015-01-12 06:09:35 UTC. The scan consists of 19 images (10 are horizontal, 8 are elevated at 45°, and 1 is overhead). The ADC counts for each pixel are plotted against zenith angle. During the scan, the camera automatically self-recalibrated by closing its internal shutter. An image taken soon after this process has been affected (Image #17, plotted in red). Affected images are characteristically less bright (lower ADC counts) than the standard images (plotted in grey). This helps to identify them. Pixels corresponding to the radio tower have been removed for clarity.

The cameras self-recalibrate using a mechanical shutter, which is positioned between the camera's sensor and the lens. The shutter resembles a black body that acts as an internal reference. Although this process regularly occurs when the cameras are operating, we have found that it typically interferes with an image taken soon after. In this dissertation, the images that have been disrupted by a camera's self-recalibration will occasionally be referred to as shutter-affected. On the other hand, normal (or standard) images are taken to mean the images unaffected by this process.

We are familiar with the shutter-affected images, as they are characteristically less bright (fewer ADC counts) than the standard images that are captured (see Figure 4.18). For our purposes of cloud detection, this can sometimes be problematic. Because we detect clouds by setting a threshold above the effective temperature of a clear sky (refer to Chapter 5), it is possible that clouds may not be identified in the less bright (or cooler) images. Typically, ~ 2 images per night are affected by a camera's internal recalibration. The process can interfere with any images during a scan, and the shutter-affected images seem to occur at random intervals of time.

Although it can be a problem for us that the affected images appear cooler, it does helps to identify them. They can, for example, easily be recognised through visual inspection of the images. Nevertheless, it can be laborious to spot them when analysing large amounts of cloud camera data. In this section, I first outline a method that can automatically identify these problematic images. Secondly, corrections for the affected images are developed.

4.4.1 Identifying problematic camera images

The Auger cloud cameras have two scanning routines. One is a 5 image panorama that covers an entire fluorescence detector (FD) field of view. The second is a 19 image full-sky mosaic. Both types of scan have a good amount of image overlap, which is particularly useful for identifying whether the data may have been affected by a camera's self-recalibration. The affected images always appear less bright than the normal images, and this characteristic helps to identify them. Through analysing images that have overlapping regions, it allows for the brightness of each image to be compared. Subsequently, the less bright (or shutter-affected) images can be determined.

Each camera pixel has an approximate 0.14° instantaneous field of view. This is too small to practically compare overlapping regions pixel-by-pixel. In order to simplify the process, a more straightforward approach is to consider larger $1^\circ \times 1^\circ$ regions of the sky. For each image, the mean pixel value (in ADC counts) within a $1^\circ \times 1^\circ$ region can be calculated. Where there are at least two images that view the same portion of the sky, a difference in the image brightness (also in ADC counts) for that region can be calculated. This can be extended to each of the $1^\circ \times 1^\circ$ regions that have overlapping images.

The results can be used to find an average difference in image brightness. Although some disagreements between images can be expected (mostly due to the warmer image edges), the overall discrepancies are small in contrast to those caused by a camera's internal recalibration.

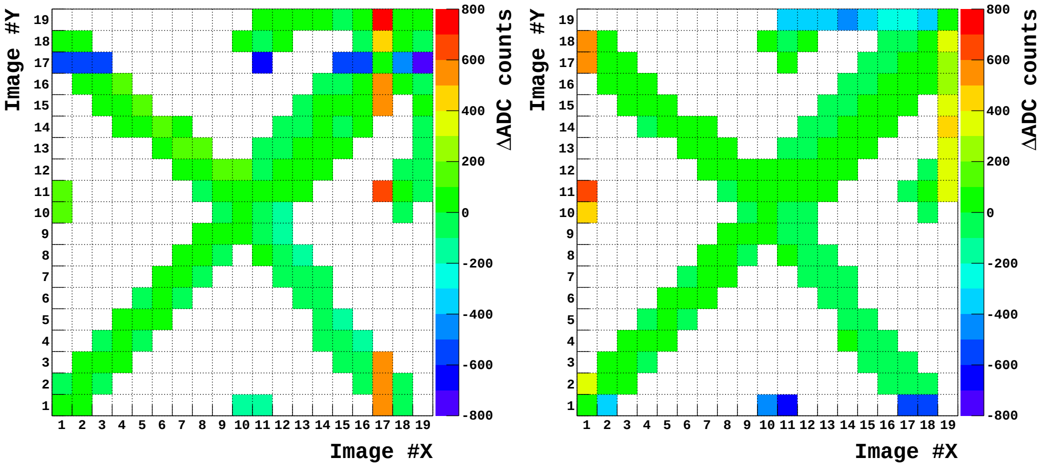

(a) **(b)**

Figure 4.19: Two examples of identifying images that have been affected by the camera's shutter during a self-recalibration process. The measurements have been made by the Los Leones infrared cloud camera. Both examples are of full-sky scans that consist of 19 individual images, with a good amount of image overlap. The colour scale gives the mean brightness difference (ΔADC) between images (Image #Y minus Image #X) in ADC counts. Only images that overlap have entries in this matrix arrangement. For instance, Image #1 has some overlap with Image #2, #10, #11, #17, and #18. Affected images are less bright than normal images, and can be determined by comparing the overlapping regions. **(a)** A scan that began at 2015-01-12 06:09:35 UTC. Image #17 was likely affected by the shutter. **(b)** A scan that began at 2015-07-23 00:57:47 UTC. It can be inferred that both Image #1 and Image #19 were affected by the shutter.

To illustrate how overlapping images can be compared, consider the two examples presented in Figure 4.19. Figure 4.19a shows brightness comparisons between one image (Image #Y) and another image (Image #X) from a full-sky scan (same as Figure 4.18). In this matrix arrangement, only the images that have some amount of overlap contain entries. The entries show the mean brightness difference (ΔADC) between images (Image #Y minus Image #X) in ADC counts. For instance, Image #1 partially overlaps with

Image #2, #10, #11, #17, and #18. It can be seen that the brightness of Image #1 is similar to most of the overlapping images, given that many of the matrix entries have ΔADC $\simeq 0$ ADC counts. In fact, the majority of images agree. An exception, however, is Image #17, which is consistently found to be less bright than the others. The disagreements can be seen graphically by the coloured strips in the figure. Consequently, Image #17 is likely to have been affected by the camera's self-recalibration. Applying the same process to the full-sky scan in Figure 4.19b, it can be inferred that Image #1 and Image #19 have been affected by the recalibration.

In both examples, the problematic images were determined graphically. To automate the detection, a threshold on the mean brightness difference can be set. I choose to identify the images when their average brightness is less than all overlapping images by at least 200 ADC counts.

The overall success at identifying images using this technique is rather high (about 92 % based on 380 full-sky scans with the Los Leones camera). It is, however, computationally intensive to search for potentially affected images in every scan performed by our cameras (about 100 scans per night for each camera). Images that are affected by a camera's self-recalibration appear to be a random occurrence, but I have typically found the procedure to impact upon ~ 2 images per night. Despite this, the majority of clouds have effective temperatures that are warm enough to stand out in the images. Most importantly, the affected images are not necessarily within the fields of view of the FDs. For these reasons, identifying the shutter-affected images is not required in most situations. Taken together, I have a successful method that can identify problematic images when necessary.

Apart from the random occurrences of self-recalibration, an update to the cloud camera software resulted in a period (May 5th, 2015 through to March 8th, 2016) where particular images were always affected by the shutter. These images were the first and last of the scanning sequence (i.e. Image #1 and Image #19 in a full-sky scan, or Image #1 and Image #5 in a FD field of view scan). Corrections that can be made to the data are discussed presently.

4.4.2 Method

As discussed in the earlier section, the automatic self-recalibration of our infrared cameras tends to disrupt the images captured soon after. In order to correct the affected images, consider how they differ from the standard images (for example, Figure 4.18). Not only do the affected images appear less bright, but their behaviour deviates from what we normally observe. In particular, there is a sudden decrease in brightness towards one of the image edges. On the other hand, normal images instead have brighter image edges due to thermal emission from the optics of our cloud camera systems. This discrepancy suggests that images affected by the internal calibration require a unique flat-field correction.

Applying the same procedure as before (see Section 4.3.1), new flat-field corrections can be obtained for the shutter-affected images. The flat-field corrections were determined using thermal radiation from the zenith for clear night skies, which can act as a uniform source of illumination. Previously, overhead cloud camera images were used because the brightness of a clear sky at 8–14 µm wavelengths remains fairly constant across those particular images (the infrared brightness varies slowly with zenith angle). In the same way, the new flat-field corrections are generated from overhead camera images. The difference now, however, is to only use images that have been disrupted by a camera's self-recalibration.

Data from the entire clear night sky catalogue (listed in Table 4.1) have been used to obtain the new flat-field corrections. In order to select only the shutter-affected measurements, an automated technique of identifying the problematic images has been used (see Section 4.4.1). Each Auger cloud camera has a unique flat-field correction that can be applied to the shutter-affected images. These new flat-field corrections have been generated using 384 images from the Los Leones camera, 387 images from the Los Morados camera, 372 images from the Loma Amarilla camera, and 402 images from the Coihueco camera.

An interesting study is to see how the new flat-field corrections (for the shutter-affected images) differ from those previously obtained (for the standard images, see Figure 4.14). However, as was found earlier, the extent of the flat-field corrections depends on the operating temperature of the camera's sensor (T_{cam}). In order to make a comparison, the new flat-field corrections have also been evaluated at a nominal operating temperature of 300 K. The results are presented in Figure 4.20.

The new flat-field corrections are considerably different to those determined in Section 4.3. What stands out in particular are the cooler image edges (also seen in Figure 4.18). As stated earlier, this is a rather unusual feature since the image edges tend to be the warmest regions of our infrared images. As we discovered with our previous corrections, the shapes of the flat-fields for the Los Leones and Loma Amarilla cameras appear to be similar, whilst those for the Coihueco and Los Morados cameras are revealed to be different. In particular, the Coihueco and Los Morados cameras show two misshapen bright regions on either side of the images. This is also surprising because the brightness is expected to increase radially from the image centre due to the increased optical path length through the cloud camera optics. However, we do not know exactly how the shutter affects the pixels' gains

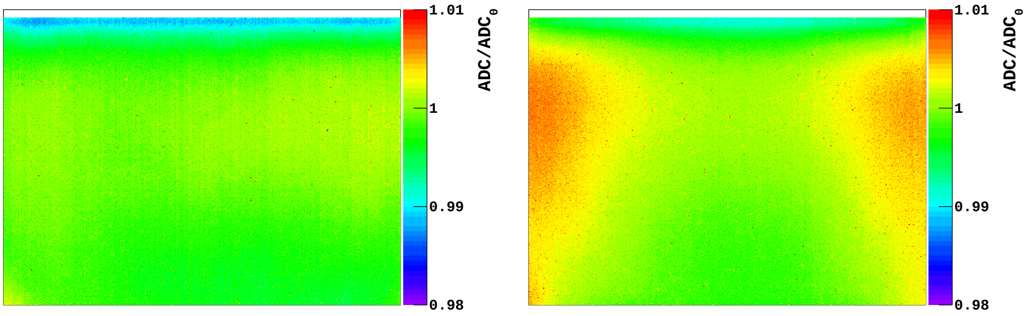

(a) Los Leones **(b)** Los Morados

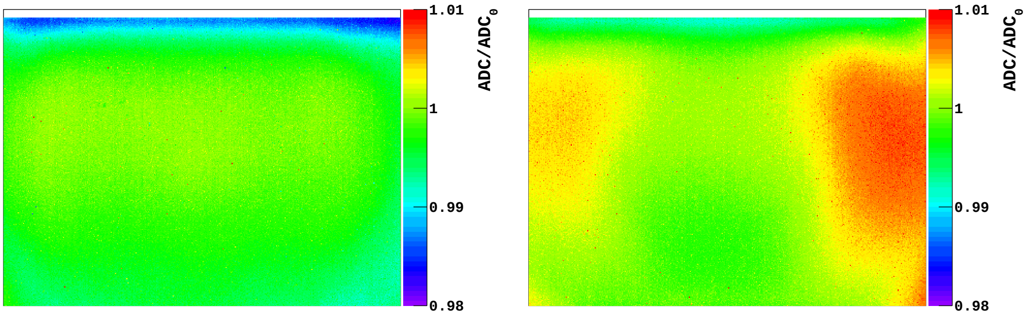

(c) Loma Amarilla **(d)** Coihueco

Figure 4.20: New flat-field correction templates for the Auger cloud cameras. These new corrections can be applied to the images affected by a camera's self-recalibration. The relative increase in brightness (ADC/ADC$_0$) of each pixel has been compared to a reference (ADC$_0$). The reference is the average value of the pixels within 2° of the image centre. Multiple clear sky images have been used to generate these templates. The extent of the corrections depends on camera sensor temperature (T_{cam}). To compare between the templates for each camera, the flat-field corrections have been adjusted to what they would be for a T_{cam} at 300 K. Comparisons can also be made to the flat-field corrections previously obtained for normal (unaffected) images (refer to Figure 4.14). Further details may be found in Section 4.3.1.

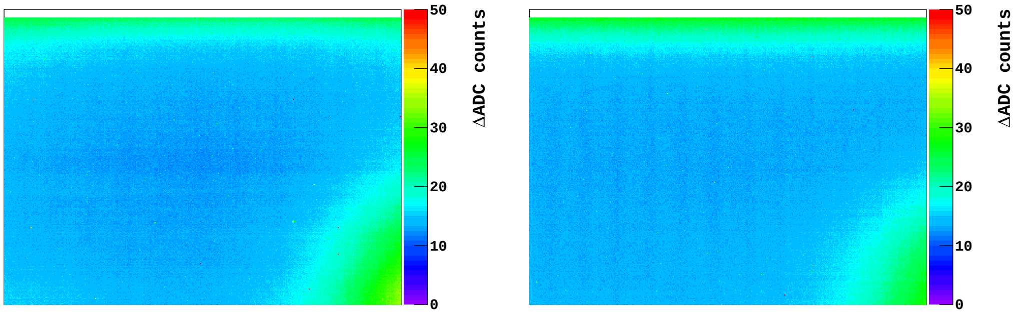

(a) Los Leones **(b)** Los Morados

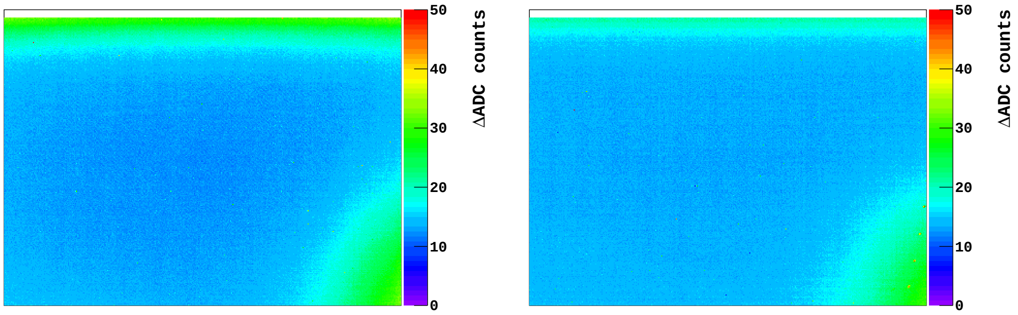

(c) Loma Amarilla **(d)** Coihueco

Figure 4.21: The root-mean-square error (RMSE) of the pixels after the new flat-field corrections have been applied to the Auger cloud camera images. These new corrections can be applied to images that have been affected by a camera's self-recalibration. Errors (in ADC counts) have been calculated by comparing the true value of a pixel to a prediction made using the new flat-field correction (see Equation 4.9 and Equation 4.10). This is repeated for a set of flat-field images. The overall error after applying the correction is ~ 20 ADC counts, which is approximately 0.2 K. This is similar to what was obtained for the normal (unaffected) images (refer to Figure 4.15). Details regarding the calculations may be found in Section 4.3.1.

during the self-recalibration process.

The performance of the new flat-field corrections can also be assessed. One way of evaluating the performance is to determine the error (in ADC counts) introduced to a pixel's value after the correction has been applied. A small error suggests an effective flat-field correction. Following the same procedure as before, the root-mean-square error (RMSE) for each pixel can be calculated using Equation 4.10. The findings are shown in Figure 4.21. Overall, the new flat-field corrections typically result in an error of 20 ADC counts, or approximately 0.2 K. Despite the stark contrast to the shapes of the previous (normal) flat-field images, the magnitudes of the errors are similar, and suitably small for our needs. Additionally, when the flat-field corrections are applied to the images, they do not appear to introduce a systematic error into the data (performing the same analysis as found in Figure 4.16).

In summary, I have two unique flat-field corrections for each Auger cloud camera. The first had been developed previously, and are used improve the quality of most images. The second are instead only used for the images disrupted by a camera's self-recalibration. To demonstrate how this works, consider applying both types of flat-field corrections to the earlier example in Figure 4.18. The images with the corrections applied are presented in Figure 4.22. A comparison between the two figures reveals the majority of image artefacts have been removed. However, what stands out now is that the shutter-affected image is offset by some amount. In order to completely correct the image, the brightness must be adjusted.

Another way of representing the data in Figure 4.22 is to group the pixel values by zenith angle. I have chosen to bin the data at 1° intervals, and the mean value within each bin is displayed by an open circle marker in the figure. Each marker is placed at the centre of a bin. The shutter-affected image (with red markers) and standard images (with black markers) have been treated separately. Pixels that correspond to the warm communications tower are omitted, for the reason that they considerably brighten the mean values, and bias the results. When the data are presented in this manner, it clarifies the adjustment in brightness that is required.

In each zenith angle bin containing both shutter-affected and standard image data, an offset ($\Delta\mathrm{ADC}_{\mathit{offset}}$) can be determined. This is calculated by subtracting the shutter-affected mean value from the standard images' mean value (i.e. the difference between a black marker and red marker at a given zenith angle). The blue arrowed line in Figure 4.22 illustrates the offset for the bin centred at 58.5°, where the standard images are found, on average, to be 442 ADC counts brighter than the shutter-affected image. Taking this concept, $\Delta\mathrm{ADC}_{\mathit{offset}}$ can be calculated for each bin. The results are presented

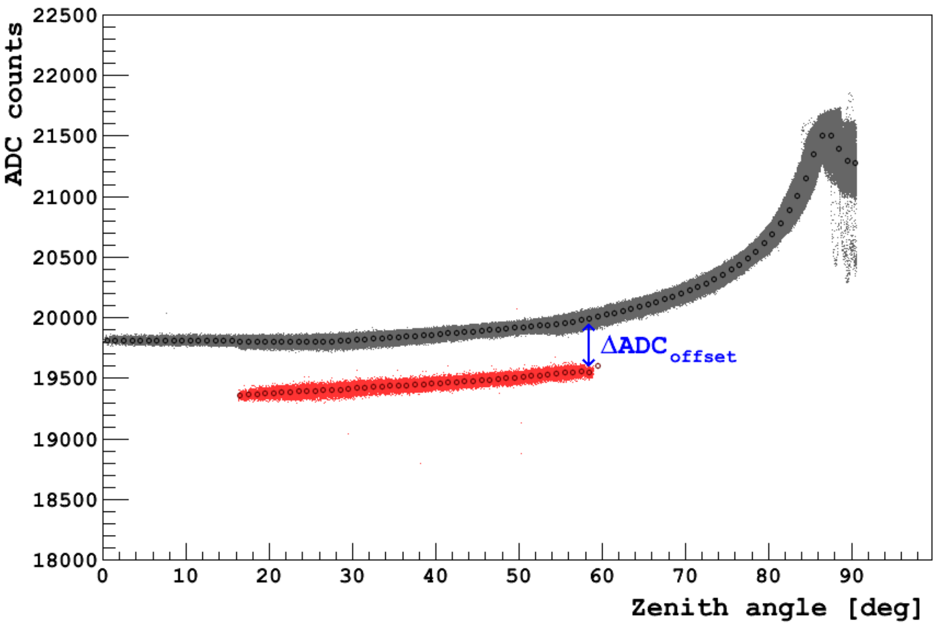

Figure 4.22: A full-sky scan captured with the Los Leones infrared cloud camera (the same as Figure 4.18). The ADC counts for each pixel are plotted against zenith angle. Image #17 (plotted in red) has been affected by the camera's shutter during an automatic self-recalibration. Two types of flat-field corrections have been applied, one for the standard images (plotted in grey), and another for the shutter-affected image. Also shown are the mean values of pixels in 1° zenith angle bins (open circle markers). These can be used to adjust the brightness of the shutter-affected image, which has been offset by some amount ($\Delta\mathrm{ADC}_{offset}$). Illustrated by the blue arrowed line is a 442 ADC count offset for the bin centred at 58.5°.

in Figure 4.23.

Only a single value is needed (per scan) to adjust the brightness of an image. The value can be determined by finding a mean brightness offset $\langle\Delta\mathrm{ADC}_{offset}\rangle$, which I calculate by averaging $\Delta\mathrm{ADC}_{offset}$ at zenith angles less than 80° (indicated by the dashed line in the figure). Using this quantity, the intensity of each pixel (with a value $\mathrm{ADC}_{affected}$) can then be adjusted to a corrected value ($\mathrm{ADC}_{corrected}$). That is:

$$\mathrm{ADC}_{corrected} = \mathrm{ADC}_{affected} + \langle\Delta\mathrm{ADC}_{offset}\rangle\,. \tag{4.11}$$

Following on with the previous example, $\langle\Delta\mathrm{ADC}_{offset}\rangle$ in Figure 4.23 is 410 ADC counts. Consequently, the image can be corrected by adding 410 ADC

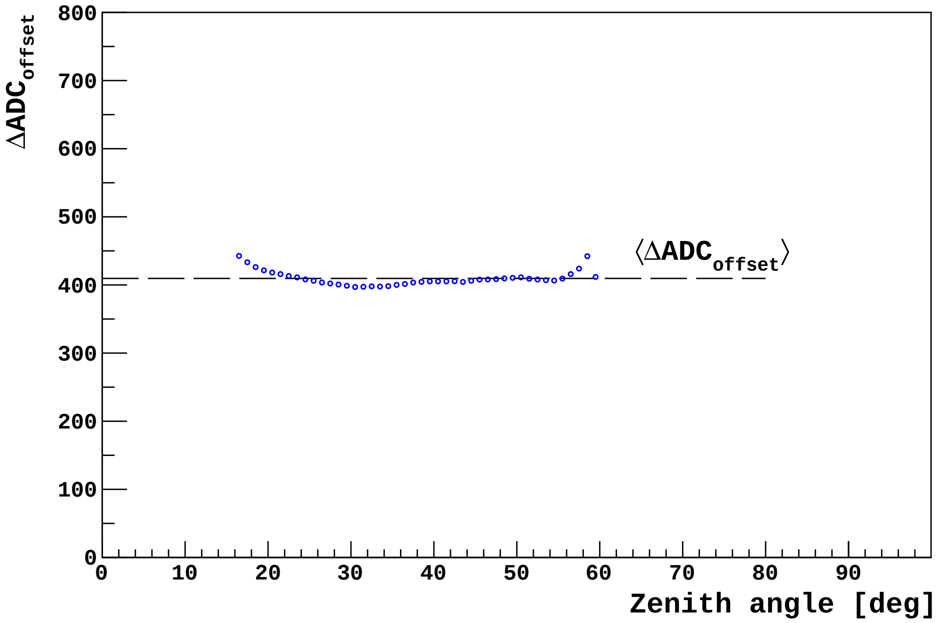

Figure 4.23: Brightness offsets at different zenith angles for an image disrupted by a camera's shutter. Data are derived from the open circle markers in Figure 4.22, which contain the mean pixel values for the shutter-affected image and standard images in fixed 1° zenith angle bins. In each bin, a $\Delta\mathrm{ADC}_{offset}$ is plotted by subtracting the shutter-affected mean value from the standard images' mean value. The mean brightness offset $\langle\Delta\mathrm{ADC}_{offset}\rangle$ for this particular example is 410 ADC counts. At this time the temperature of the camera's sensor (T_{cam}) was 316.51 K.

counts to all shutter-affected pixel values. Although this modification works well for this particular scan (where the camera sensor temperature is 316.51 K), the main goal is to find corrections that can be applied to any cloud camera measurement.

One quantity that has been found to affect the brightness corrections is the temperature given by the camera's sensor (T_{cam}). Given that T_{cam} has been found to influence the pixel output and temperature resolution of our cameras (see Section 4.1.2), some relationship with the sensor temperature can be expected. In addition, the problem we are trying to solve is due to the mechanical shutter. Since the shutter is at a similar temperature to the camera's sensor, it seems reasonable that our corrections would also depend on T_{cam} in some way.

In order investigate the relationship between the brightness corrections

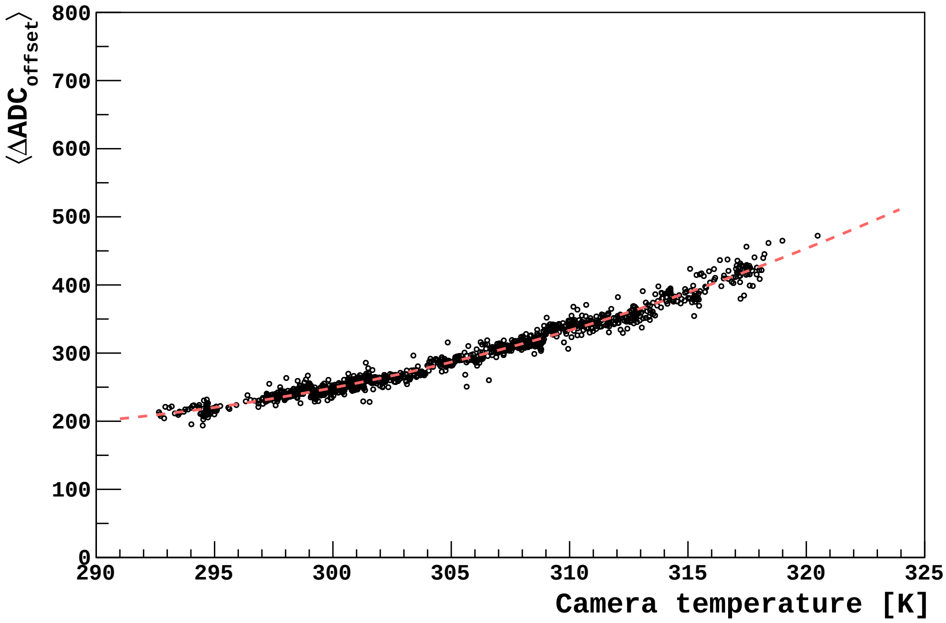

Figure 4.24: The mean brightness offset $\langle\Delta\mathrm{ADC}_{offset}\rangle$ for Los Leones cloud camera images disrupted by the camera's self-recalibration. Data are plotted against the temperature of the camera's sensor (T_{cam}). Results have been obtained from multiple scans, and show that the offsets depend on T_{cam}. The relationship can be parameterised using a parabolic function (see Equation 4.12). A fit to the data is indicated by the dashed red parabola, and the fit parameters are listed in Table 4.2.

and sensor temperature, cloud camera measurements from the clear night sky catalogue have been used (listed in Table 4.1). Given that the data generally covers most of a night, and are throughout various seasons, the camera sensor temperatures span a wide range of operating values. When the new flat-field corrections were developed, I was restricted to using full-sky scans with overhead shutter-affected images. However, data from both types of scanning routines can be used for this analysis, provided that at least one shutter-affected image has been identified. Hence, this larger dataset consists of 1197 scans from the Los Leones camera, 1191 scans from the Los Morados camera, 1162 scans from the Loma Amarilla camera, and 1226 scans from the Coihueco camera.

Figure 4.24 shows the mean brightness offset for the Los Leones cloud camera plotted against the camera sensor temperature. It is apparent from the figure that the brightness corrections depend on the camera sensor temperature, with the warmer sensor temperatures requiring a greater correction. In fact, the same trend is revealed for each of the Auger cloud cameras (see Figure 4.25). A simple function that can describe the behaviour reasonably well is a parabola. The brightness correction can then be estimated from:

$$\langle \Delta \mathrm{ADC}_{\mathit{offset}} \rangle = a_2 \times T_{cam}^2 + a_1 \times T_{cam} + a_0 \,. \tag{4.12}$$

Since we already monitor the temperature of a camera's sensor, the remaining information needed to correct an image are the fit parameters a_2, a_1, and a_0. The parameters are unique to each Auger camera, and the values listed in Table 4.2.

Site	$\boldsymbol{a_2}$	$\boldsymbol{a_1}$	$\boldsymbol{a_0}$
Los Leones	0.1789	-100.70	14356
Los Morados	0.2573	-147.55	21377
Loma Amarilla	0.2309	-129.82	18456
Coihueco	0.3064	-175.78	25454

Table 4.2: Parameters used to adjust the brightness of measurements disrupted by a camera's self-recalibration (see Equation 4.12). Each Auger cloud camera has a separate adjustment. The parameters have been determined from the fitted data in Figure 4.25.

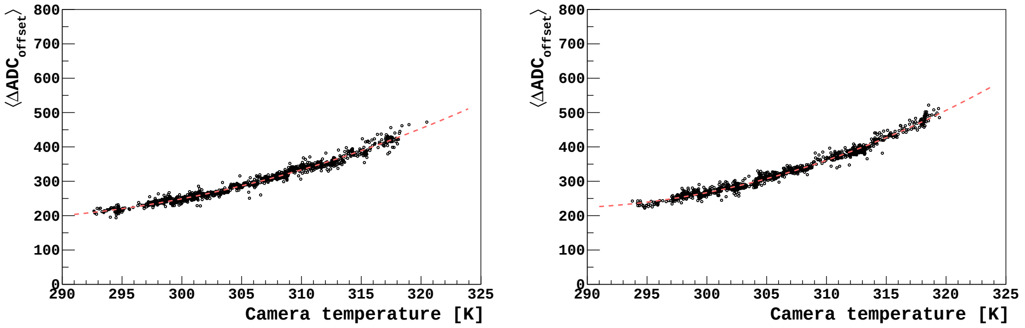

(a) Los Leones **(b)** Los Morados

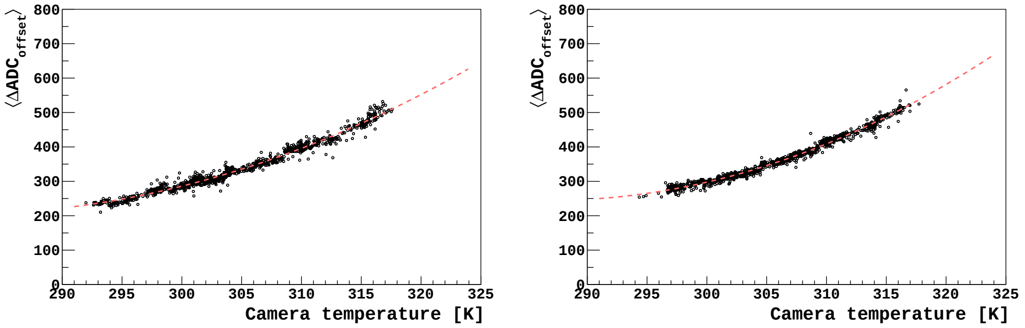

(c) Loma Amarilla **(d)** Coihueco

Figure 4.25: The mean brightness offset $\langle \Delta\mathrm{ADC}_{offset} \rangle$ for images affected by a camera's self-recalibration plotted against the temperature of the camera's sensor (T_{cam}). Results are for each of the Auger cloud cameras, and have been obtained using data from multiple scans. This shows that the brightness offset depends on T_{cam}. For each camera, the dependence can be parameterised using a parabolic function (see Equation 4.12). The fits are indicated by the dashed red parabolas, and the fit parameters are listed in Table 4.2.

In summary, the techniques presented in this section can successfully correct the images that have been disrupted by our cameras' self-recalibration. The first step is to apply a (unique) flat-field correction to the affected image, which removes the majority of the artefacts. However, the overall brightness of the image is too low, and must then be adjusted. The brightness can be corrected using Equation 4.11, where the extent of the correction $\langle \Delta \mathrm{ADC}_{offset} \rangle$ depends on the camera sensor temperature (T_{cam}). This can be estimated using Equation 4.12, and the values listed in Table 4.2. An example of cloud camera data with these image corrections applied is shown in Figure 4.26.

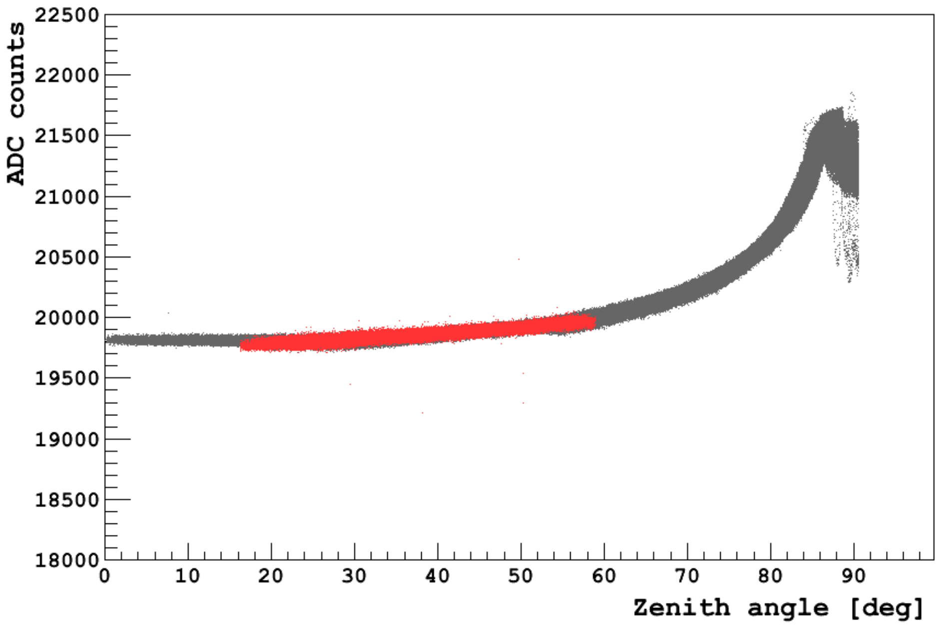

Figure 4.26: A full-sky scan captured with the Los Leones infrared cloud camera (the same as Figure 4.18). The ADC counts for each pixel are plotted against zenith angle. Image #17 (plotted in red) has been affected by the camera's shutter during an automatic self-recalibration. Separate flat-field corrections have been applied to the shutter-affected image and standard images (plotted in grey). The brightness of the shutter-affected image has been corrected using Equation 4.11, Equation 4.12, and the parameters listed in Table 4.2.

4.5 Temperature calibration

Our thermal imaging cameras convert incident radiation into a measurable signal, which is digitised, and expressed in ADC counts. However, the signal is also affected by the properties and settings of the camera. Hence, the camera's digital output is not an absolute measurement. In order to analyse the images quantitatively, it is necessary to calibrate our instruments. I calibrate our cameras to give temperature readings such that measurements only depend on the infrared brightness within their fields of view.

In addition, there are many clouds that resemble black bodies at infrared wavelengths (see Section 2.4). As a result, an accurate temperature reading would be similar to the actual temperature of the cloud. On the other hand, clear skies usually appear much cooler because the atmosphere is transparent at those wavelengths. Only near the horizon is the atmosphere opaque, which is due to the increased path lengths of water vapour and carbon dioxide (the main emitting gases within our cameras' band-pass). Since optically thick (opaque) clouds generally appear warmer than their surroundings, they can often be detected at infrared wavelengths by setting some temperature threshold above a clear sky background. One major benefit to applying a temperature threshold rather than an ADC count threshold is that it removes any dependence on a camera's properties and settings. For these reasons, we wished to determine temperature calibrations for the Auger cloud cameras.

Pixel outputs can be converted from ADC counts into temperatures whenever a quantitative analysis is required. The technique that I have developed has built upon my understanding of the thermal cameras from my previous studies (see Section 4.1). In particular, it was found that a camera's digital output (in ADC counts) of an object (at a temperature T_{obj} in Kelvin) in the field of view can be converted into a temperature reading using the following equation:

$$\mathrm{ADC} = m\left(T_{cam}\right) \times T_{obj} + C\left(T_{cam}\right) \ , \tag{4.2}$$

where the slope $m(T_{cam})$ and offset $C(T_{cam})$ have some dependence on the temperature of the camera's sensor (T_{cam} in Kelvin). Finding temperature conversions for the Auger cloud cameras by completing Equation 4.2 is the focus of this section.

4.5.1 Method

When a temperature calibration was produced for one of our cameras at the University of Adelaide, I had access to an experimental black body. However, the Auger cameras had to be calibrated separately because they were already installed at the Observatory. Although we cannot directly calibrate our cameras in a controlled environment, we do have access to data from single-pixel radiometers that also operate at the Observatory (see Section 3.2). The infrared radiometers had been calibrated (using a black body) in a laboratory, and have a somewhat similar spectral response to our thermal imaging cameras.

There are two single-pixel radiometers that are installed at the Pierre Auger Observatory. One is located at the Central Laser Facility (CLF), and another at the eXtreme Laser Facility (XLF). Both instruments are mounted in order to record an overhead sky temperature within the band-pass of their detectors. In addition, the local temperature of the detector canister is recorded by a thermistor. Because the detector canisters are fixed close to the ground, the thermistor temperatures approximate the local air temperature at surface-level. The data are collected at 5 minute intervals, and I have chosen to calibrate our cameras using sky temperature and thermistor temperature measurements from the CLF radiometer.

Although we do not have access to an experimental black body, the atmosphere appears optically thick when viewing angles near the horizon. This is because as the atmospheric path length is extended, the intervening matter is increased. In particular, this applies to the greenhouse gases, which are well-known absorbers at these wavelengths (see Section 2.2.2). Since the atmosphere is opaque near the horizon, it resembles a black body. The temperature of the black body is similar to the local air temperature, which is recorded by our radiometer's thermistor. In fact, if the radiometer was inclined such that sky temperatures near the horizon were measured instead, the observations would then closely approximate the thermistor temperatures [53]. Consequently, a thermistor temperature can be used as a proxy for the sky temperature at the horizon. On the other hand, the overhead sky temperature of a clear night sky (which is not optically thick) is much lower. Our radiometers typically measure sky temperatures that are about 20 K cooler than their thermistor temperatures. We know from preliminary measurements that our camera's response to scene temperatures is approximately linear over this temperature range (see Figure 4.3). Taken together, we can perform two-point temperature calibrations using the radiometer data. One point is the temperature of a clear sky measured at the zenith, and the second point is the temperature of the thermistor (a proxy for the sky temperature at the horizon). I have chosen to use clear skies rather than overcast skies because it maximises the temperature difference between both points, and will reduce errors in the calibrations.

The actual procedure for calibrating the Auger cloud cameras is analogous to what we had before (see Section 4.1.2). A camera's digital output

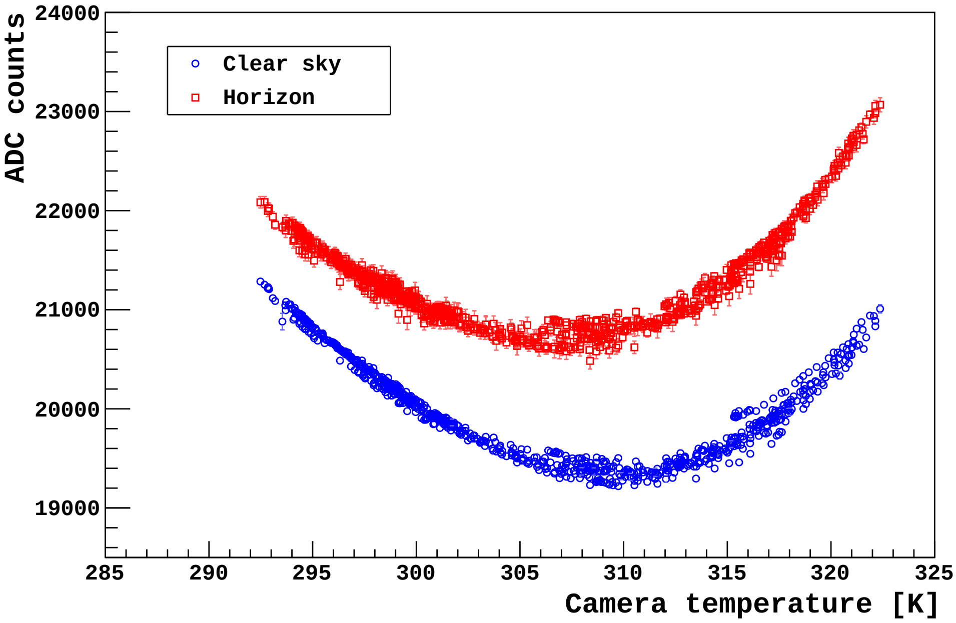

Figure 4.27: Measurements made by the Los Leones cloud camera at different camera sensor temperatures (T_{cam}). Data were collected from full-sky scans performed on clear nights. The camera viewed two types of scenes (at different effective temperatures). One is a cool scene of a clear sky near the zenith (blue circle markers). The second is a warmer scene near the horizon (red square markers). Mean pixel values for each scene are plotted, and errors derived from the standard deviation of those averaged values. The relationship between the camera's digital output (in ADC counts) and T_{cam} resembles preliminary measurements that were carried out in a laboratory (see Figure 4.4). Unlike the preliminary experiment, these scene temperatures are not constant, which leads to some scatter in the ADC counts (for example, at a given T_{cam}, a cooler night sky will be less bright and have fewer counts).

(in ADC counts) corresponding to known temperatures in the field of view can be used to produce a calibration curve. In this case, an overhead measurement made with a camera can be compared to the radiometer's sky temperature. Likewise, a measurement taken near the horizon with a camera can be compared to the radiometer's thermistor temperature. Consequently, only full-sky scans are used for this particular analysis. To illustrate the similarities to the previous method, consider the data presented in Figure 4.27. In this example, clear sky and horizon measurements captured with the Los Leones camera have been plotted against the temperature of the camera's sensor (see Appendix C, Figure C.1 for the other cameras). What is interesting about this figure is the parabolic-like response of the camera with T_{cam}. These findings resemble my earlier results in the laboratory (see Figure 4.4), and demonstrate why we wish to compensate for the camera's temperature. Following a similar procedure to my earlier experiments, the influence that T_{cam} has on the calibrations can be understood by performing two-point temperature calibrations at different camera temperatures. In order to span a wide range of operating temperatures, I use measurements that were recorded on various clear nights throughout the years.

The data are selected from a catalogue of clear night skies that I have constructed, which may be found listed in Table 4.1. This calibration method relies on the CLF radiometer and cloud camera both measuring a clear sky. Hence, to increase the likelihood that a clear sky was also observed at the CLF, there is an additional requirement that at least 2 operating cloud cameras are viewing a clear sky (also listed in the table). Each of the Auger cameras has a unique temperature calibration. The calibrations have been produced using data from 847 full-sky scans captured with the Los Leones camera, 743 full-sky scans captured with the Los Morados camera, 672 full-sky scans captured with the Loma Amarilla camera, and 469 full-sky scans captured with the Coihueco camera. I apply a flat-field correction to each image in a scanning sequence in order to remove various artefacts (see Section 4.3). In addition, data that may have been affected by a camera's self-recalibration are identified and corrected (see Section 4.4). To see how the temperature calibrations are produced, first consider how to calibrate the Los Leones cloud camera.

Figure 4.28 shows Los Leones cloud camera data (with the flat-field corrections) that was collected on a clear night from a full-sky scan. The camera's signal can be calibrated by making comparisons to contemporaneous temperature measurements recorded with the CLF radiometer. Camera data near the zenith corresponds to the radiometer's sky temperature. I have chosen the mean value of all pixels at zenith angles less than 20° to represent a camera's sky measurement (about 11500 pixels), which has been shaded blue in the figure. The error in that measurement is taken as the standard deviation of those pixels' values. On the other hand, camera data close to the horizon corresponds to the radiometer's thermistor temperature. In practice, it was less straightforward to define a camera's horizon measurement. One of the reasons is because the brightness increases rapidly with zenith angle towards the ground (as the air mass increases). In addition, the brightness of the ground tends to decrease when our cameras view the Earth's surface at an angle (possibly due to a decrease in the surface emissivity at increased

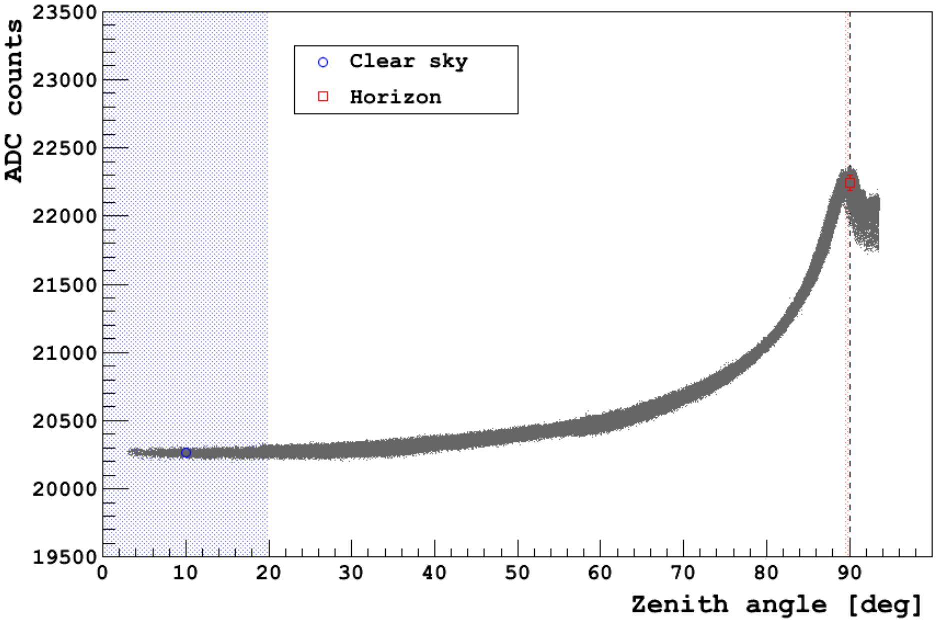

Figure 4.28: Measurements made by the Los Leones infrared cloud camera when the night sky was clear. The full-sky scan began at 2015-02-11 01:51:49 UTC, and consists of 19 images (10 are horizontal, 8 are elevated at 45°, and 1 is overhead). A flat-field correction has been applied to the images. The ADC counts for each pixel are plotted against zenith angle (θ). A sky measurement of 20266 ± 17 ADC counts (blue circle marker) is taken as the average of all pixels at $\theta < 20°$ (shaded blue). A horizon measurement of 22244 ± 58 ADC counts (red square marker) is taken as the average of all pixels at $89.5° < \theta < 90°$. A temperature calibration can be determined using both measurements, as the corresponding temperatures are known from a single-pixel radiometer operating at the Observatory. The horizon is indicated by the dashed black line. Only pixels that have an azimuth (ϕ) within $0° < \phi < 80°$ are used.

viewing angles [93]). Consequently, a horizon measurement can be strongly influenced by the pointing accuracy of the pixels. I shift the pointing directions so that the sky brightness peaks at 90° (indicated by the dashed line in the figure), which seems reasonable since it is where the atmosphere is thickest. Through visual inspection, the maximum brightness for a clear sky aligns well with the horizon (note that some pointing directions for the Coihueco camera have an additional correction, which is discussed in Appendix B). After the adjustment, I have chosen the mean value of all pixels with zenith angles between 89.5° and 90° to represent a horizon measurement (about 3000 pixels), which has been shaded red in the figure. The error in that measurement is taken as the standard deviation of those pixels' values. In order to select pixels corresponding to a clear sky, data containing the communication towers, weather stations, and mountains in the field of view are omitted. I only use pixels within a limited azimuth range for this study, which is $0° < \phi < 80°$ for the Los Leones camera, $80° < \phi < 180°$ for the Los Morados camera, $-130° < \phi < -10°$ for the Loma Amarilla camera, and $-100° < \phi < 10°$ for the Coihueco camera.

In this particular example, the camera's sky measurement was 20266 ± 17 ADC counts (blue circle marker), and the horizon measurement was 22244 ± 58 ADC counts (red square marker). These can act as two points for the calibration provided the corresponding temperatures are known from the radiometer. All radiometer data within 7 minutes of a camera's full-sky scan (typically 3 measurements) are used to determine an average sky temperature, and an average thermistor temperature. In this case, the sky temperature was $266.2 \pm 1.2\,\mathrm{K}$, and the thermistor temperature was $290.9 \pm 0.2\,\mathrm{K}$. Figure 4.29a shows an example of a two-point temperature calibration using this information. It follows that for this particular observation, a two-point temperature calibration (given by Equation 4.2) would have a slope $m = 80.1 \pm 4.6$ ADC counts per K, and an offset $C = -1052 \pm 1351$ ADC counts. Note that the errors for m and C are correlated (see Figure 4.29b and Figure 4.29c), and the resultant uncertainty in the calibration will be discussed later (Equation 4.16 and Figure 4.31). This particular calibration curve is indicated by the dashed line in the figure. However, this was only for a single measurement when T_{cam} was $319.3\,\mathrm{K}$. The same procedure can be applied to the remaining 846 observations. Subsequently, a relationship between the two-point calibrations and the temperatures of the sensor can be explored.

Figure 4.29b and Figure 4.29c demonstrate that the two-point temperature calibrations depend on the operating temperature of the camera's sensor (see Appendix C, Figure C.2 and Figure C.3 for the other Auger cameras). These findings resemble my results from a similar experiment, which had been carried out in a laboratory (refer to Figure 4.5b and Figure 4.5c). One likeness is how the calibration slopes appear steeper as T_{cam} becomes warmer. A variable slope (which relates to a temperature difference) implies that the temperature resolution of our camera changes, which could also be seen in Figure 4.27 since the separation between clear sky (at cool temperatures) and horizon measurements (at warmer temperatures) would depend on T_{cam}. Whilst the slopes (Figure 4.29b) tend to determine the separation between those measurements, the offsets (Figure 4.29c) can explain the parabolic-like features observed in Figure 4.27. Since we monitor T_{cam}, temperatures in the

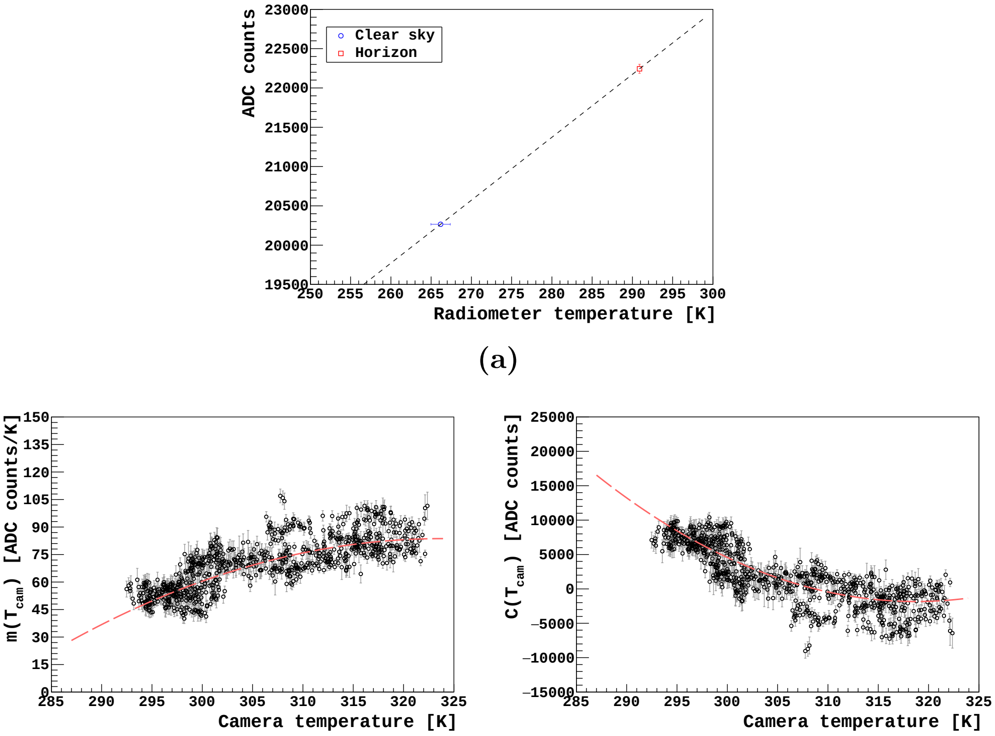

(a)

(b) **(c)**

Figure 4.29: **(a)** An example of a two-point temperature calibration by comparing Los Leones cloud camera data (the same full-sky scan as Figure 4.28) to single-pixel radiometer temperature measurements. One point is an overhead clear sky temperature at 266.2 ± 1.2 K (blue circle marker). The second point is a thermistor temperature at 290.9 ± 0.2 K (red square marker), which can act as a proxy for the sky temperature near the horizon. On average, the camera's digital output was 20266 ± 17 ADC counts for the first point, and 22244 ± 58 ADC counts for the second point. Consequently, this particular two-point calibration (dashed line) would have a slope $m = 80.1 \pm 4.6$ ADC counts per K, and an offset $C = -1052 \pm 1351$ ADC counts. Note the errors for m and C are correlated. At this time the temperature of the camera's sensor (T_{cam}) was 319.3 K. **(b)** Various two-point calibration slopes $m(T_{cam})$ calculated from 847 full-sky scans captured on clear nights with the Los Leones camera at different sensor temperatures. A fit to the data is indicated by the dashed red parabola, and the fit parameters listed in Table 4.3. **(c)** Various two-point calibration offsets $C(T_{cam})$ calculated from 847 full-sky scans captured on clear nights with the Los Leones camera at different sensor temperatures. A fit to the data is indicated by the dashed red parabola, and the fit parameters listed in Table 4.4. Errors derived from uncertainties in the temperature and camera measurements.

field of view can be estimated provided that both a slope $m(T_{cam})$ and an offset $C(T_{cam})$ are known (see Equation 4.2). I parameterise the relationship between $m(T_{cam})$ and $C(T_{cam})$ with the temperature of the camera's sensor (T_{cam} in Kelvin) using parabolic functions. That is:

$$m\left(T_{cam}\right) = m_2 \times T_{cam}^2 + m_1 \times T_{cam} + m_0\,, \tag{4.13}$$

and

$$C\left(T_{cam}\right) = C_2 \times T_{cam}^2 + C_1 \times T_{cam} + C_0\,. \tag{4.14}$$

Table 4.3 lists the slope parameters m_2, m_1, and m_0 for each of the Auger cameras, and the offset parameters C_2, C_1, and C_0 are listed in Table 4.4. It follows that a camera's signal can be converted into a temperature measurement by estimating a value for $m(T_{cam})$ and $C(T_{cam})$ at a known sensor temperature. That is, the temperature of an object in the field of view (T_{obj} in Kelvin) can be calculated from:

$$T_{obj} = \frac{\mathrm{ADC} - C\left(T_{cam}\right)}{m\left(T_{cam}\right)}\,, \tag{4.15}$$

by substituting in Equation 4.13 and Equation 4.14.

Site	$\boldsymbol{m_2}$	$\boldsymbol{m_1}$	$\boldsymbol{m_0}$
Los Leones	-0.0410166	26.5623	-4216.77
Los Morados	-0.0275126	17.8670	-2849.20
Loma Amarilla	-0.0454892	29.0644	-4552.77
Coihueco	-0.0157238	10.1617	-1603.07

Table 4.3: Parameters used to estimate a slope m for a temperature calibration (see Equation 4.13), which can be used to convert a camera's digital output into a temperature reading (see Equation 4.15). The parameters have been determined from the fitted data in Appendix C, Figure C.2. Note that several decimal places have been listed, as an estimated temperature reading can be sensitive to these parameters.

To investigate the success of the camera's temperature calibration, comparisons can be made to the radiometer measurements. Since the camera has been calibrated using the radiometer data, such a comparison should ideally yield a one-to-one relationship. Figure 4.30a shows the calibrated

Site	C_2	C_1	C_0
Los Leones	18.2398	-11627.60	1851260
Los Morados	14.4149	-9159.16	1461280
Loma Amarilla	18.4040	-11631.50	1835040
Coihueco	12.5747	-7846.99	1235800

Table 4.4: Parameters used to estimate an offset C for a temperature calibration (see Equation 4.14), which can be used to convert a camera's digital output into a temperature reading (see Equation 4.15). The parameters have been determined from the fitted data in Appendix C, Figure C.3. Note that several decimal places have been listed, as an estimated temperature reading can be sensitive to these parameters.

camera data (using Equation 4.15) plotted against the radiometer temperatures. The good correlation between both quantities suggests that the camera has been calibrated reasonably well. Another way to test the accuracy of the temperature readings is by looking at distributions of differences between the predicted and known temperatures. Figure 4.30b shows distributions of the calibrated camera data minus radiometer temperatures for the clear sky and horizon measurements. The overall distribution (clear sky and horizon measurements) has a mean of 0.5 K and a standard deviation of 2.4 K, which also implies the temperature calibration performs rather well.

Although this particular temperature calibration appears to work reasonably well, there are some rather unusual features in Figure 4.30a. Namely, the disagreements that seem to occur when the radiometer's sky temperature (at $\sim 250\,\mathrm{K}$) and thermistor temperature (at $\sim 268\,\mathrm{K}$) are cooler. In order to determine the source of these discrepancies, consider the data presented in Figure 4.31 (see Appendix C, Figure C.4 for the other Auger cameras). The figure shows the temperature differences (camera minus radiometer) plotted against the temperature of the camera's sensor, and shows that the accuracy of the temperature calibration varies with T_{cam}. In particular, the calibration appears to be least accurate at lower sensor temperatures. This was also when the radiometer's sky temperatures and horizon temperatures were cooler, which explains the discrepancies seen in Figure 4.30a. The findings suggest that T_{cam} may not be correctly compensated for by parameterising the calibration slopes and offsets using Equation 4.13 and Equation 4.14. Despite this, the data implies that a modification can be made to the tem-

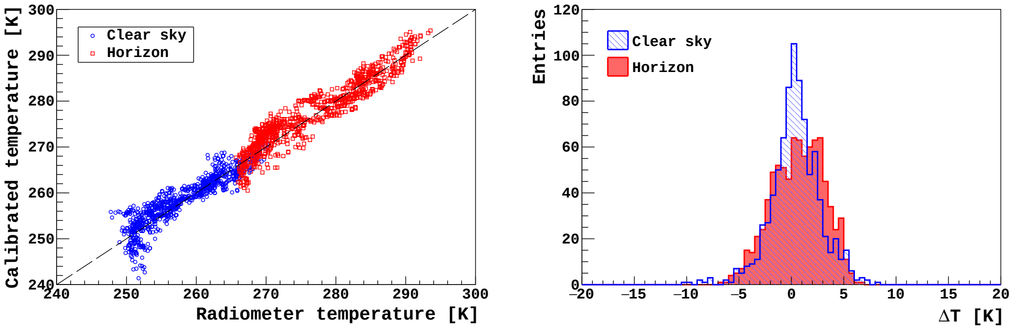

(a) (b)

Figure 4.30: **(a)** Measurements made with the Los Leones cloud camera after a temperature calibration (a prediction based on Equation 4.15) has been applied. The data are plotted against temperature measurements recorded by a single-pixel infrared radiometer, which has a somewhat similar spectral response to the camera, and provided a truth that was used to calibrate the instrument. Camera data near the zenith can be compared to the radiometer's sky temperature (blue circle markers), and camera data near the horizon (red square markers) can be compared to the radiometer's thermistor temperature (a proxy for the sky temperature near the horizon). A one-to-one relationship is indicated by the dashed line. **(b)** Distributions of the temperature differences using the same data. Camera measurements (prediction) minus the radiometer measurements (assumed truth) are shown. Overall, the combined temperature difference distribution has a mean of 0.5 K and a standard deviation of 2.4 K.

perature calibration in order to account for these discrepancies. That is:

$$T_{obj} - T_{true} = r\left(T_{cam}\right) \,, \tag{4.16}$$

where T_{obj} (in Kelvin) is the predicted temperature of a field of view object (using Equation 4.15), T_{true} (in Kelvin) is the true temperature measurement (known from the radiometer in this case), and $r(T_{cam})$ is a correction term that can improve the accuracy of the temperature reading. I parameterise $r(T_{cam})$ using a cubic function, which is given by:

$$r\left(T_{cam}\right) = r_3 \times T_{cam}^3 + r_2 \times T_{cam}^2 + r_1 \times T_{cam} + r_0 \,, \tag{4.17}$$

where the parameters r_3, r_2, r_1, and r_0 for each of the Auger cameras are listed in Table 4.5.

In summary, a camera's digital output (in ADC counts) can be converted into a temperature reading (T_{obj} in Kelvin) provided that the temperature

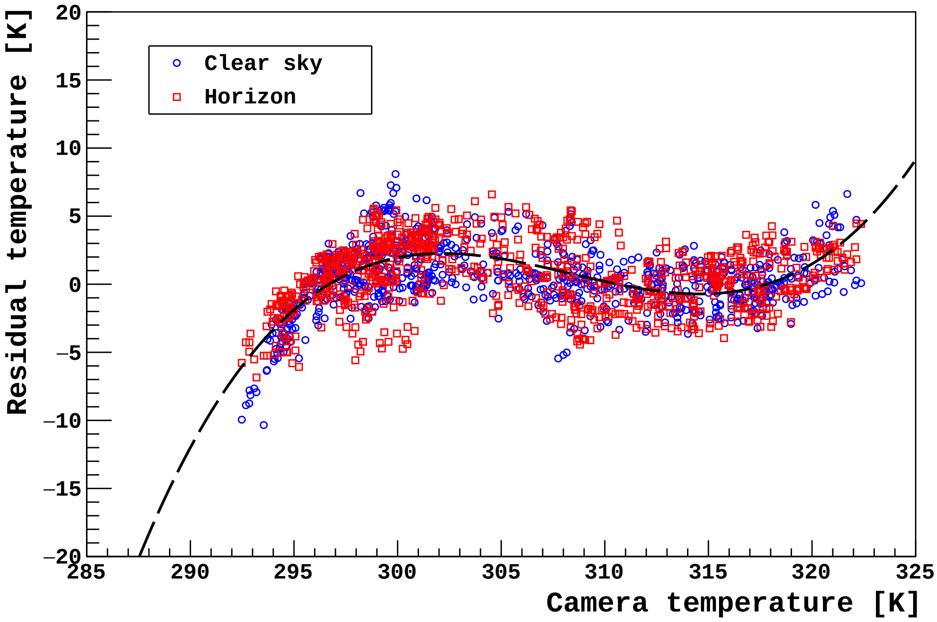

Figure 4.31: Discrepancies in temperature measurements (using calibrated data) made with the Los Leones cloud camera plotted against the temperature of the camera's sensor (T_{cam}). The residual temperature is the calibrated temperature using camera data (a prediction based on Equation 4.15) minus the temperature measurement made by a single-pixel radiometer (the assumed truth). Camera data near the zenith can be compared to the radiometer's sky temperature (blue circle markers), and camera data near the horizon (red square markers) can be compared to the radiometer's thermistor temperature (a proxy for the sky temperature near the horizon). The results show the residuals depend on T_{cam}, and suggests that a modification can be made to the temperature calibration. The dependence can be parameterised using a cubic function (see Equation 4.17). A fit to the data is indicated by the dashed black curve, and the fit parameters are listed in Table 4.5.

of the camera's sensor (T_{cam} in Kelvin) is known. That is:

$$T_{obj} = \frac{\mathrm{ADC} - C\left(T_{cam}\right)}{m\left(T_{cam}\right)} - r\left(T_{cam}\right) \ , \tag{4.18}$$

where $m(T_{cam})$ can be substituted with Equation 4.13, $C(T_{cam})$ with Equation 4.14, and $r(T_{cam})$ with Equation 4.17.

The improvements to the temperature calibration can be seen by comparing the updated calibrated camera measurements to the radiometer temper-

Site	r_3	r_2	r_1	r_0
Los Leones	0.00312422	-2.89030	890.940	-91507
Los Morados	0.00482160	-4.46037	1374.950	-141233
Loma Amarilla	0.00279310	-2.56030	782.061	-79604
Coihueco	0.00816265	-7.52799	2313.410	-236892

Table 4.5: Parameters used to estimate a correction term r (see Equation 4.17) for a more accurate temperature calibration (see Equation 4.16). The parameters have been determined from the fitted data in Appendix C, Figure C.4. Note that several decimal places have been listed, as the estimated temperature reading can be sensitive to these parameters.

ature measurements. Figure 4.32a shows the calibrated camera data (using Equation 4.18) plotted against the radiometer temperatures. What stands out is a stronger correlation between both quantities, with a clearer one-to-one relationship, particularly at cooler sensor temperatures. Similarly, Figure 4.32b shows new distributions of the calibrated camera data minus radiometer temperatures for the clear sky and horizon measurements. The overall distribution (clear sky and horizon measurements) has a mean of 0.0 K and a standard deviation of 1.9 K. The reduced spread in the figure also suggests an improvement to the temperature calibration.

This method of determining a temperature calibration for the Los Leones cloud camera can be successfully applied to the three remaining Auger cameras, which are installed at the Los Morados, Loma Amarilla, and Coihueco fluorescence detector sites. The intermediate steps may be found in Appendix C. Figure 4.33 shows the calibrated camera data plotted against the radiometer temperature data for each of the cameras. A good correlation between both quantities can be seen for each of the cameras, which highlights the reproducibility of the method.

In order to assess the performance of each camera's temperature calibration, it is useful to evaluate the error (in Kelvin) associated with the temperature readings. One way to quantify this error is to calculate the root-mean-square error (RMSE) by comparing the predicted temperatures (T_{pred}) to the true temperatures (T_{true}). The predicted temperatures are calculated using the temperature calibration (see Equation 4.18), and the (assumed) true temperatures are known from the radiometer. A small error corresponds to a more accurate temperature calibration, and the statistic

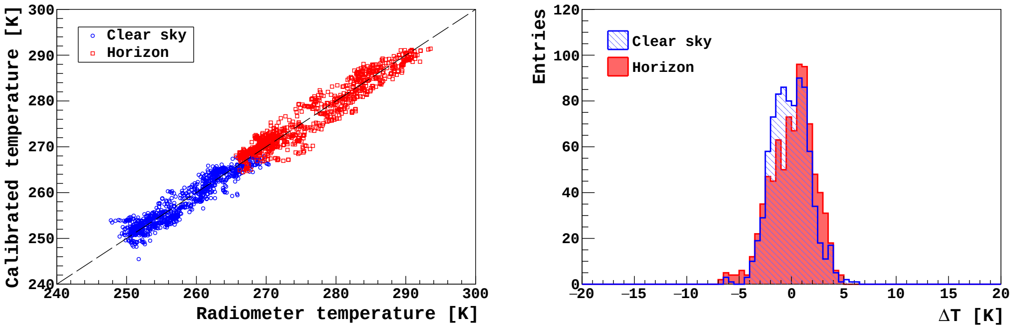

(a) **(b)**

Figure 4.32: **(a)** Measurements made with the Los Leones cloud camera after an improved temperature calibration (a prediction instead based on Equation 4.18) has been applied. The data are plotted against temperature measurements recorded by a single-pixel infrared radiometer, which has a somewhat similar spectral response to the camera, and provided a truth that was used to calibrate the instrument. Camera data near the zenith can be compared to the radiometer's sky temperature (blue circle markers), and camera data near the horizon (red square markers) can be compared to the radiometer's thermistor temperature (a proxy for the sky temperature near the horizon). A one-to-one relationship is indicated by the dashed line. **(b)** Distributions of the temperature differences using the same data. Camera measurements (prediction) minus the radiometer measurements (assumed truth) are shown. Overall, the combined temperature difference distribution has a mean of 0.0 K and a standard deviation of 1.9 K. Both figures indicate that temperature readings are, in general, more accurate (compared to Figure 4.30), particularly at cooler camera sensor temperatures.

can be calculated from the following expression:

$$\mathrm{RMSE} = \sqrt{\frac{\sum_{i=1}^{N} \left(\mathrm{T}_{pred} - \mathrm{T}_{true}\right)^2}{N}}\,, \tag{4.19}$$

where N is the total number of observations (i.e. $N = 1694$ for the Los Leones camera, which consists of 847 overhead sky measurements and 847 horizon measurements).

The RMSE is about 1.9 K for the Los Leones camera, 2.1 K for the Los Morados camera, 1.8 K for the Loma Amarilla camera, and 2.5 K for the Coihueco camera. Interestingly, the temperature calibration appears to be least accurate for the Coihueco camera. This seems reasonable considering

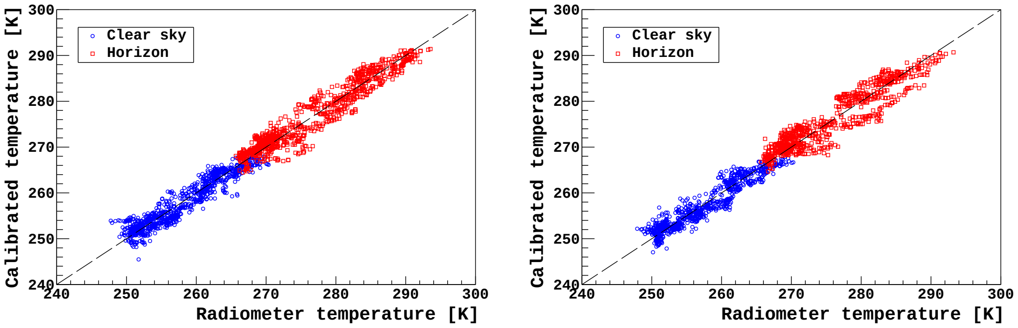

(a) Los Leones **(b)** Los Morados

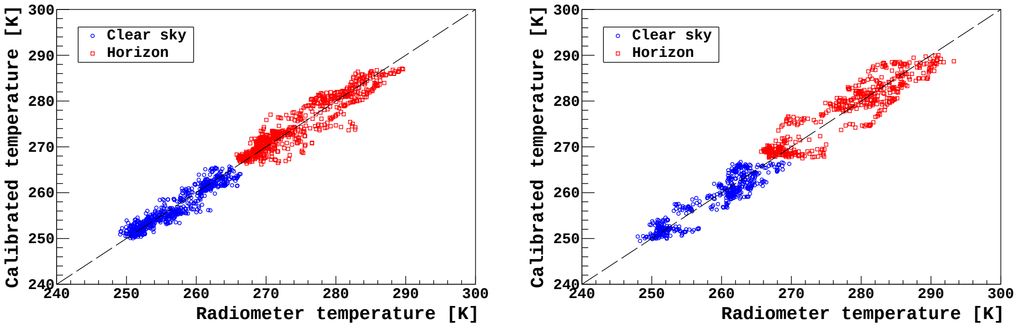

(c) Loma Amarilla **(d)** Coihueco

Figure 4.33: Measurements made with each of the Auger cloud cameras after temperature calibrations (see Equation 4.18) have been applied. The data are plotted against temperature measurements recorded by a single-pixel infrared radiometer, which has a somewhat similar spectral response to our cameras, and provided a truth that was used to calibrate the cameras. Camera data near the zenith can be compared to the radiometer's sky temperature (blue circle markers), and camera data near the horizon (red square markers) can be compared to the radiometer's thermistor temperature (a proxy for the sky temperature near the horizon). A one-to-one relationship is indicated by the dashed line in each figure. The root-mean-square error (RMSE) of the temperature calibrations (see Equation 4.19) is about 1.9 K for the Los Leones camera, 2.1 K for the Los Morados camera, 1.8 K for the Loma Amarilla camera, and 2.5 K for the Coihueco camera.

the Coihueco site is at a higher altitude (1719 m a.s.l.) than the Central Laser Facility (1401 m a.s.l.) where the radiometer is located. As a result, the Coihueco camera is expected to measure cooler sky temperatures than the radiometer, particularly towards the horizon where the temperature is similar to that of the local air temperature. Another interesting aspect of this study has been that the accuracy of the temperature calibration does not appear to be influenced by humidity. This implies that our cloud camera systems and single-pixel radiometers are equally sensitive to atmospheric water vapour, and supports our initial assumption that the cameras and radiometers share a similar spectral response.

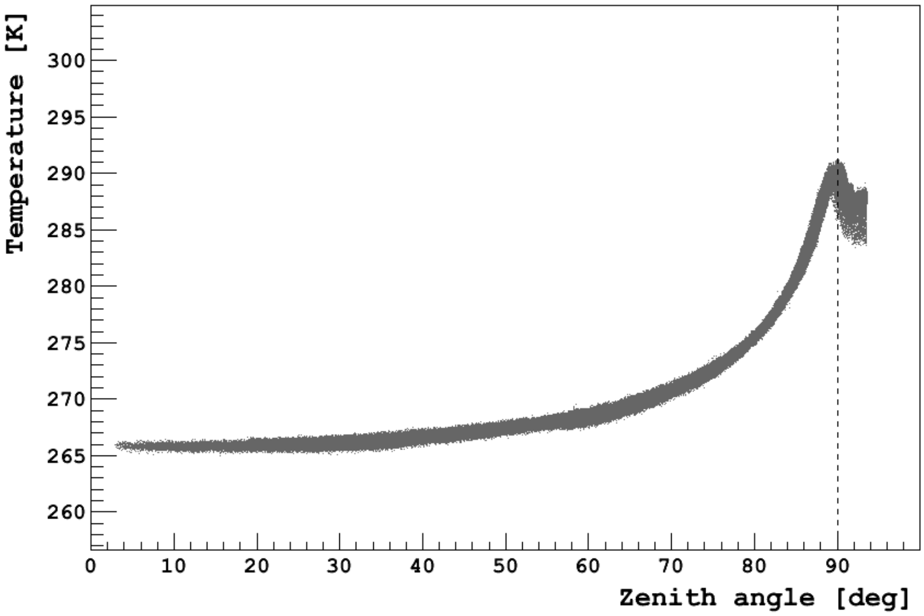

Figure 4.34: Measurements made by the Los Leones infrared cloud camera (the same full-sky scan as Figure 4.28). A flat-field correction has been applied to the images. The ADC counts for each pixel have been converted into a temperature reading (using a temperature calibration, see Equation 4.18) and plotted against zenith angle. The sky temperatures agree with contemporaneous measurements collected by a single-pixel radiometer at the Observatory, which measured the overhead sky temperature to be 266.2 K, and the thermistor temperature (a proxy for the sky temperature at the horizon) to be 290.9 K. The radiometer has a similar spectral response to the camera, and was used to calibrate the instrument.

Using the results of this study, Figure 4.34 shows an example of applying a temperature calibration to data collected with the Los Leones camera (the uncalibrated data had been shown in Figure 4.28). The sky temperature near the zenith, and towards the horizon are in good agreement with contemporaneous measurements recorded with the CLF radiometer. What is particularly useful about the data presented in this manner, is that the camera data is an absolute measurement, and can be analysed quantitatively. Analysing the cloud camera data after applying the temperature calibration, and image corrections is one of the aims in Chapter 5.

4.6 Summary

The major focus of this chapter has been to develop routines that can be used to calibrate our thermal imaging cameras, and improve the quality of our digital images. In particular, data collected by the infrared cameras operating as cloud monitors at the Pierre Auger Observatory. One aspect of my work has been to determine flat-field corrections for our cameras (see Section 4.3). The corrections not only remove image artefacts associated with the optics of our cloud camera systems, but can also account for variations in the pixel-to-pixel sensitivity. In addition, our cameras automatically perform their own corrections by self-recalibrating using a mechanical shutter. This occurs when the cameras are operating, however, we have found that it typically disrupts the next image in the sequence. Correcting for the images that have been affected by this process was the aim of Section 4.4. Finally, temperature calibrations were produced for each of the cameras (see Section 4.5). A temperature calibration allows us to convert a camera's signal into a temperature reading, which allows us to analyse the data quantitatively. The method I use to calibrate the Auger cameras had built upon my understanding of the cameras from several experiments that I had carried out (see Section 4.1), and also incorporated the image processing techniques already mentioned.

The starting point for the majority of these routines was to construct a catalogue of clear night skies (listed in Table 4.1). It would be possible to monitor the performance of these procedures in the future by adding more clear night data to the catalogue. In addition, the routines outlined in this chapter could perhaps be applied to other thermal imaging cameras, and benefit other researchers in this field.

A general outline of how I process the cloud camera data is listed below:

1. Apply flat-field corrections to the images in order to remove various artefacts (see Section 4.3.1).

2. Identify whether an image has been disrupted by the camera's mechanical shutter. An affected image requires a separate flat-field correction (see Section 4.4.2). In addition, the image brightness must also be adjusted (see Equation 4.12).
3. Convert the camera's digital output (in ADC counts) into a temperature reading (see Equation 4.18).

One of the main goals in Chapter 5 is to analyse the cloud camera data quantitatively, which occurs after these data processing steps have been applied.

Chapter 5

Cloud Mask Production

The infrared cameras operating at the Pierre Auger Observatory are used for night-time cloud detection. Since clouds can affect the reconstruction of extensive air showers using the fluorescence technique (see Section 1.2.3), the specific objective of our cameras is to identify clouds within the Auger fluorescence detectors' (FDs') fields of view. Although clouds can often be identified through visual inspection of the thermal images, our cameras collect large amounts of data (about 1000 individual images per night for each camera), and it is necessary to automate the process. The process is used to produce cloud "masks", which provide the cloud cover contained within each of the FD pixels (see Section 5.3). This information is stored into a database where they can be accessed by other members of the collaboration.

Developing an algorithm capable of identifying clouds in our images has been an important goal of my studies, and is the theme of this chapter. A key aspect to detecting clouds, however, is to distinguish the signal from a clear sky which also emits thermal radiation.

5.1 Thermal radiation from clear night skies

Most of the radiation that reaches the ground at night is thermal radiation emitted in the atmosphere (see Chapter 2). On a clear night, atmospheric emission (and absorption) is minimised within the 8–14 µm waveband. Because the vertical atmosphere has a high transmittance at these wavelengths (see Figure 2.2), the apparent emissivity of the atmosphere is rather low (about 0.2–0.6 [94]). Consequently, when our cameras are directed towards the zenith, they are detecting infrared flux contributions from multiple atmospheric layers, as well as a background radiance from outer space (about 3 K). The strength of the atmospheric contribution depends on the temperature of the gas, which depends on its height in the atmosphere [29].

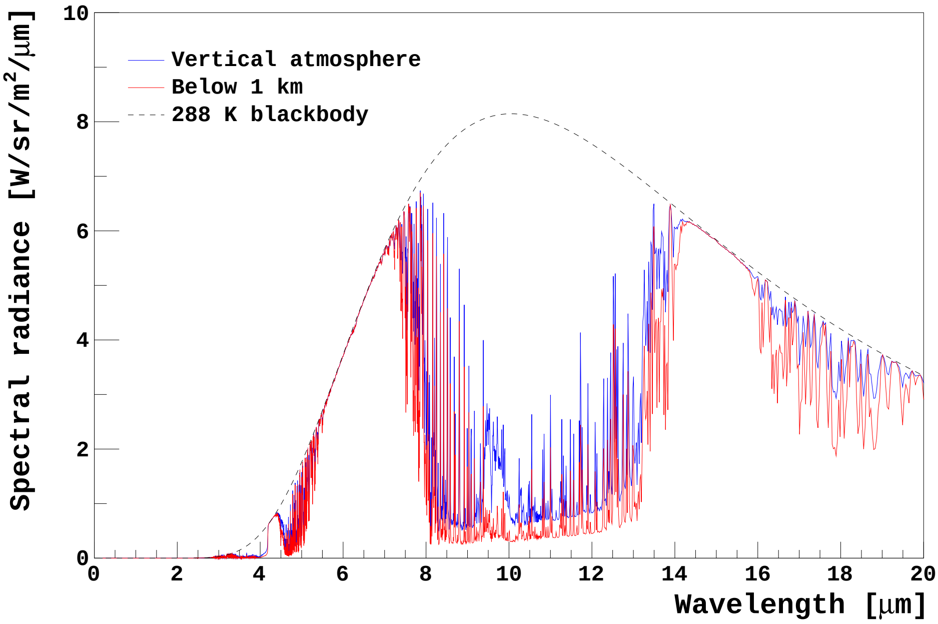

Figure 5.1: Atmospheric radiance from a clear sky reaching the ground. Simulated along a vertical track by the author with MODTRAN for a 1976 U.S. Standard Atmosphere. The emission spectrum of the entire vertical atmosphere (sea-level to space), and the emission spectrum of the lowest kilometre of atmosphere are shown. The radiance (integrated from 0–50 µm) due to the entire vertical atmosphere is $\sim 83\,\mathrm{W\,sr^{-1}\,m^{-2}}$, whereas the radiance due to the lowest kilometre is $\sim 75\,\mathrm{W\,sr^{-1}\,m^{-2}}$ (about 90 % of the entire vertical atmosphere). The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference.

Since most of the infrared flux received at ground is from the lowest kilometre of the atmosphere [32] (and illustrated by the simulation in Figure 5.1), a simple way to express the downwelling radiant flux (F_{sky} in $\mathrm{W\,m^{-2}}$) can be given by:

$$F_{sky} = \epsilon_{sky} \sigma T_{air}^4 \, , \tag{5.1}$$

where σ is the Stefan-Boltzmann constant, ϵ_{sky} is the apparent emissivity of the atmosphere (considering the entire atmosphere as a grey body), and T_{air} is the local air temperature in Kelvin. Consequently, the effective sky temperature (T_{sky} in Kelvin) may be expressed as:

$$F_{sky} = \sigma T_{sky}^4 \, , \tag{5.2}$$

which is what we observe with our thermal imaging cameras after a temperature calibration has been applied to the data (see Section 4.5).

There are various formulas for estimating ϵ_{sky} (for example [95][96][97]), which often relate the atmospheric emissivity to humidity and the ambient temperature. The formulas, however, depend on the spectral response of the detector and the geographical location, and are typically developed using long-term averages [98]. Hence, the formulas cannot be expected to describe all situations well.

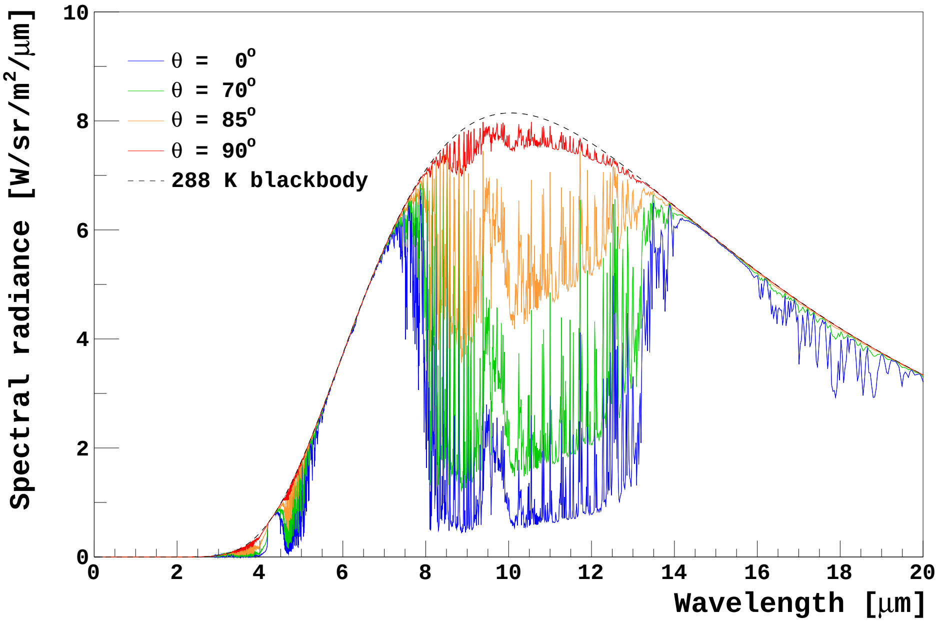

Figure 5.2: Atmospheric radiance from a clear sky reaching the ground. Simulated by the author with MODTRAN for a 1976 U.S. Standard Atmosphere for trajectories through the atmosphere at different zenith angles (θ). The infrared brightness increases with zenith angle due to the increased path lengths of water vapour and carbon dioxide, which are the main emitting gases at these wavelengths. The temperature of the Earth's surface is 288 K in this model, and the radiance of a black body at that temperature is shown for reference. Note that at $\theta = 90°$ the atmosphere is not optically thick in this simulation due to the curvature of the Earth in the model.

The emissivity of the atmosphere (or effective sky temperature, see Equation 5.2) also varies with zenith angle (for example [53][94][99][100]). This is due to the increasing optical depths of water vapour and carbon dioxide, which are the main emitting gases at long-wave infrared wavelengths (see Section 2.2.2). Consequently, the atmospheric infrared flux increases from the zenith towards the horizon. This is illustrated in Figure 5.2, and shows the clear sky brightness for a ground-based observer, which I have simulated at different zenith angles. What particularly stands out is the rapid increase in brightness with zenith angle within the 8–14 µm waveband. Because our cloud cameras are sensitive to radiation at those wavelengths, it is difficult to distinguish a cloud's thermal emission from a clear sky background when viewing angles towards the horizon (since the atmosphere is opaque and bright). The variations in brightness with zenith angle can also be seen in Figure 5.3a, which presents the effective sky temperature (which relates to the infrared brightness) on a clear night as a function of zenith angle, as measured with the Los Leones cloud camera.

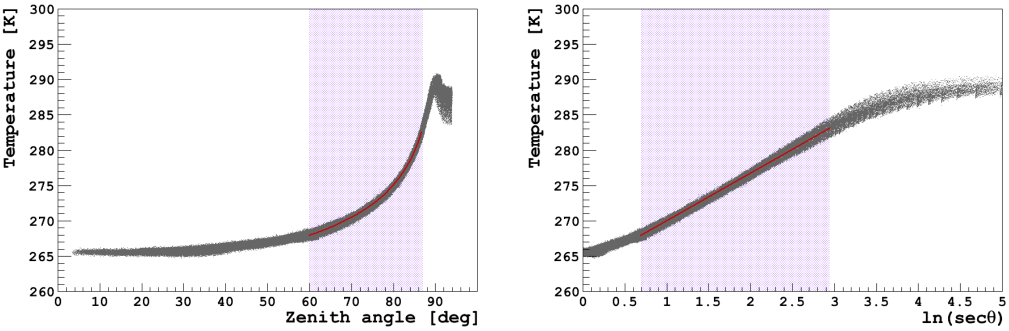

(a)

(b)

Figure 5.3: Measurements made by the Los Leones infrared cloud camera. The full-sky scan began at 2015-02-11 01:36:49 UTC, and consists of 19 images (10 are horizontal, 8 are elevated at 45°, and 1 is overhead). A flat-field correction has been applied to the images, and a temperature calibration has been applied to the data. A temperature reading for each pixel is plotted against **(a)** zenith angle θ, and **(b)** $\ln(\sec\theta)$. Equation 5.3 can describe the change in sky temperature with zenith angle rather well between $60° < \theta < 87°$ (shaded purple in both figures), which is the approximate field of view of the Auger fluorescence detectors. A fit to the data in that region is indicated by the red curve (and red line). For this particular observation, $A = 263.3\,\mathrm{K}$ and $B = 6.7\,\mathrm{K}$. Note that as $\theta \to 90°$ then $\ln(\sec\theta) \to \infty$.

A simple equation that can describe the change in effective sky temperature T_{sky} with zenith angle θ had been determined in [53], and given as:

$$T_{sky}(\theta) = A + B \times \ln(\sec\theta) \ , \tag{5.3}$$

where A and B are fitted parameters. The empirical relationship had been derived from measurements collected by single-pixel infrared radiometers (one was the same model discussed in Section 3.2). I have found this particular equation to describe our cloud camera data reasonably well. This is not surprising given that the radiometer used in the original study had a similar spectral response to our cameras, and the same model (the radiometer located at the Central Laser Facility) had also been used to calibrate our instruments in Section 4.5.

One way to test the applicability of Equation 5.3 to our cloud camera data is to plot the sky temperatures against $\ln(\sec\theta)$. A linear relationship between both quantities would suggest the model accurately describes the temperature variations with zenith angle. It can be seen in Figure 5.3b that the relationship is approximately linear, however, there are some discrepancies at smaller, and larger zenith angles. Despite this, I have found the function generally fits our data between $60° < \theta < 87°$ (approximately $0.7 < \ln(\sec\theta) < 3$) rather well. The fitted region has been shaded in both figures, and a fit to the data indicated by the red curve (or red line in Figure 5.3b). We are particularly interested in data within this region because it is the approximate field of view of the Auger fluorescence detectors. When the data are presented such as in Figure 5.3b, it clarifies the A parameter as the offset, and the B parameter as the slope of the line. For this particular observation, $A = 263.3\,\mathrm{K}$ (about $2.5\,\mathrm{K}$ cooler than the measured sky temperature at the zenith) and $B = 6.7\,\mathrm{K}$. Determining parameters that can be be used to predict clear sky temperatures for other observations will be discussed presently.

5.1.1 Parameterising effective sky temperatures with zenith angle

In order to estimate the clear sky temperature at a given zenith angle, values for A and B (the parameters in Equation 5.3) are required. In this section, I relate the parameters to physical quantities that we can measure.

I parameterise the model using cloud camera data recorded on clear nights. The data are selected from a catalogue of times that I have constructed, which have been listed in Table 4.1. Since the data are fitted over a limited range of zenith angles rather than the entire sky, measurements from both scanning routines (full-sky and the FD field of view) can be used for this particular analysis. In order to remove various artefacts, a flat-field correction is applied to each image within a scanning sequence (see Section 4.3). In addition, data that may have been affected by a camera's self-recalibration are identified and corrected (see Section 4.4). The effective sky temperatures used for this study are obtained by converting the digital output of each pixel into a temperature reading (see Section 4.5). This study uses data from 2971 scans captured with the Los Leones camera, 2474 scans captured with the Los Morados camera, 1749 scans captured with the Loma Amarilla camera, and 1691 scans captured with the Coihueco camera.

For each camera scan, a fit to the data yields the parameters A and B (for example, see Figure 5.3). Measurable quantities that have been found to affect these parameters are the air temperature and atmospheric water vapour. Since most of the downwelling radiation from the sky is emitted low in the atmosphere, the ambient ground temperature can be expected to influence the results (see Equation 5.1). In addition, our cameras are sensitive to atmospheric water vapour, and it seems reasonable that humidity would also impact on our measurements.

The local air temperature can be determined using our cloud cameras. Because the atmosphere is opaque near the horizon, it resembles a black body, and the temperature of the body is similar to the ambient temperature. As a result, when our cameras are viewing the horizon, the recorded sky temperatures approximate the temperature of the surrounding air. I define the air temperatures measured with our cameras as the average temperature reading of pixels with zenith angles between $89.5^\circ < \theta < 90^\circ$ in an unobstructed region of the sky (about 3000 pixels), and the details on this horizon measurement may be found in Section 4.5.1. The error in the air temperature is taken as the standard deviation of those pixels' values and is typically about 1 K. There is an additional uncertainty of about 2 K due to the accuracy of our temperature calibration.

There are several ways of quantifying the amount of moisture in the atmosphere (see Section 2.3). Given that our cameras survey the night sky, they are not only sensitive to the humidity at surface-level, but they also receive infrared flux contributions from water vapour in multiple atmospheric layers. Hence, the most useful quantity to use in this study is the total amount of water vapour contained within an atmospheric column. This can be expressed as the total precipitable water vapour (TPWV), which is the depth of water in an atmospheric column, if all the water precipitated as rain. I calculate the TPWV using height-dependent atmospheric profiles from the Global Data Assimilation System (see GDAS in Section 3.6). The details regarding the calculation may be found in Section 2.3.1. Since the GDAS profiles are only available every 3-hours (and our cameras scan at 5 minute intervals), I estimate the total precipitable water vapour at the time of a camera's scan from a linear interpolation to the data. Typical values of TPWV in these datasets range from 2–17 mm.

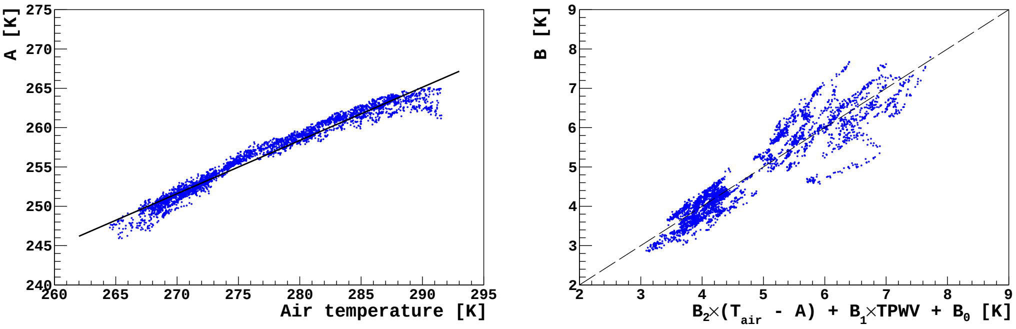

(a) **(b)**

Figure 5.4: **(a)** Fit parameter A (in Equation 5.3) plotted against the air temperature for measurements made with the Los Leones camera. The air temperature can be determined with the camera, and taken as the sky temperature near the horizon. The relationship has been parameterised using a linear function (see Equation 5.4). A fit to the data is indicated by the black line, and the fit parameters listed in Table 5.1. **(b)** Fit parameter B (in Equation 5.3) plotted against predicted values of B using the air temperature (T_{air}), the parameter A, and the total precipitable water vapour (TPWV), which is the depth of water vapour in an atmospheric column, if all the water precipitated as rain. Predictions have been made using Equation 5.5 and the parameters listed in Table 5.2. A one-to-one relationship is indicated by the dashed line.

Figure 5.4a shows the fit parameter A plotted against different air temperatures using measurements made with the Los Leones camera. I parameterise the relationship with a linear function, which is indicated by the black line fitted to the data. Consequently, the parameter A in Equation 5.3 can be estimated from the air temperature (T_{air} in Kelvin) using the following:

$$A = A_1 \times T_{air} + A_0\,, \tag{5.4}$$

where the parameters A_1 and A_0 for each of the Auger cameras are listed in Table 5.1. The general performance of the model can be seen in Figure 5.6a, which shows a histogram of predicted A values (using the model) minus the true values (from fits to the data) for all four Auger cameras. The overall uncertainty in A can be estimated from the distribution, and is found to be about 0.9 K.

Site	$\boldsymbol{A_1}$	$\boldsymbol{A_0}$ [K]
Los Leones	0.676	69.0
Los Morados	0.760	45.9
Loma Amarilla	0.688	66.3
Coihueco	0.651	76.6

Table 5.1: Parameters used to estimate A (see Equation 5.4) for each of the Auger cloud cameras, which can then be used to predict the effective temperature of a clear night sky (see Equation 5.3). The parameters have been determined from the fitted data in Figure 5.7 (at the end of this section). Overall, the predicted values of A typically have an error of about 0.9 K (see Figure 5.6a).

The remaining information needed to estimate clear sky temperatures comes from the fit parameter B in Equation 5.3. The way that B (which is the slope of the line in Figure 5.3b) has been defined relates the parameter to the temperature difference between $\theta = 60°$ and $\theta = 87°$. On cloudy nights, however, it is not always possible to determine B directly from measurements since clouds appear warmer than a clear sky. A quantity that can be used instead is $(T_{air} - A)$, which somewhat reflects this temperature difference. This is because the change in sky temperature (temperature at the horizon minus the assumed temperature at the zenith) with zenith angle (90° minus 0°) will be the same assuming the linear relationship in Equation 5.3. One benefit to using this quantity is that it only relies on the air temperature. Another advantage is that temperatures measured near the horizon tend to approximate the air temperature regardless of whether clouds are present (since the atmosphere is opaque). It was found empirically that B was also influenced by the amount of atmospheric moisture. A formula that I have found that can be used to estimate the parameter B is given by the following:

$$B = B_2 \times (T_{air} - A) + B_1 \times TPWV + B_0 \,, \tag{5.5}$$

where T_{air} is the air temperature in Kelvin, A can be estimated using Equation 5.4 and is also in Kelvin, and $TPWV$ is the total precipitable water vapour in millimetres. Fit parameters B_2, B_1, and B_0 for each of the Auger cameras are listed in Table 5.2.

Figure 5.4b shows the fit parameter B plotted against predictions based on Equation 5.5. A one-to-one relationship is also shown for reference, and

Site	$\boldsymbol{B_2}$	$\boldsymbol{B_1}$ [K mm^{-1}]	$\boldsymbol{B_0}$ [K]
Los Leones	0.233	0.15	-1.1
Los Morados	0.284	0.15	-2.0
Loma Amarilla	0.253	0.15	-2.0
Coihueco	0.181	0.19	-0.4

Table 5.2: Parameters used to estimate B (see Equation 5.5) for each of the Auger cloud cameras, which can then be used to predict the effective temperature of a clear night sky (see Equation 5.3). The parameters have been determined from fitted data. Estimates using Equation 5.5 and the values in this table are compared to the true fit parameters in Figure 5.8 (at the end of this section). Overall, the predicted values of B typically have an error of about 0.4 K (see Figure 5.6b).

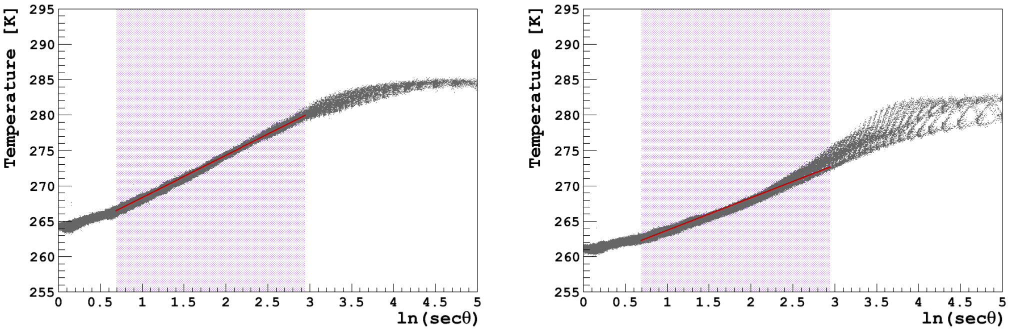

(a) **(b)**

Figure 5.5: Measurements made by the Los Leones infrared cloud camera on two different clear nights. **(a)** A full-sky scan that began at 2015-12-20 03:41:07 UTC. **(b)** A full-sky scan that began at 2015-12-11 05:44:44 UTC. A temperature reading for each pixel has been plotted against $\ln(\sec\theta)$, where θ is the zenith angle. Equation 5.3 has been fitted to the data between $60^\circ < \theta < 87^\circ$ (approximately $0.7 < \ln(\sec\theta) < 3$) in order to describe the change in sky temperature with zenith angle. Fits to the data are indicated by the red lines, and the fitted region has been shaded purple in both figures. The empirical formula better describes the measurements in (a). Note that as $\theta \to 90^\circ$ then $\ln(\sec\theta) \to \infty$.

it can be seen there are some discrepancies in the model. The disagreements, however, were not found to correlate with humidity, air temperature, or the temperature of the camera's sensor. Despite this, when inspecting the fitted data (such as in Figure 5.3), it is clear that Equation 5.3 fits the data better in some cases than in other cases (for example, see Figure 5.5). I suspect this to be the main cause of the discrepancies observed. The performance of the fit, however, does seem to depend on the day, which leads to some of the streaking effects seen in the figure. The overall performance of the model can be seen in Figure 5.6b, which shows a histogram of predicted B values (using the model) minus the true values (from fits to the data) for all four Auger cameras. The uncertainty in B can be estimated from the distribution, and is found to be about 0.4 K.

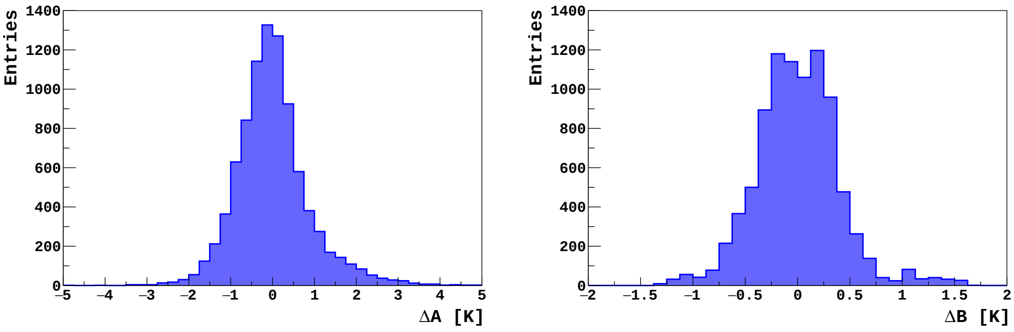

(a) (b)

Figure 5.6: **(a)** A distribution showing predicted A values (in Equation 5.3) minus true values obtained from fits to the data. Predictions have been made using an empirical model (Equation 5.4). Similar distributions can be produced for each of the Auger cloud cameras, and the combined distribution (for all four cameras) has been shown. Overall, the combined difference distribution has a mean of 0.0 K and a standard deviation of 0.9 K. **(b)** A distribution showing predicted B values (in Equation 5.3) minus true values obtained from fits to the data. Predictions have been made using an empirical model (Equation 5.5). Similar distributions can be produced for each of the Auger cloud cameras, and the combined distribution (for all four cameras) has been shown. Overall, the combined difference distribution has a mean of 0.0 K and a standard deviation of 0.4 K.

In summary, clear night sky temperatures vary with zenith angle, and can be described reasonably well for our cloud camera data between $60^\circ < \theta < 87^\circ$ with Equation 5.3. The equation has two parameters, A and B, which I relate to the air temperature and the atmospheric water vapour content.

Empirical formulas for estimating A and B are given by Equation 5.4 and Equation 5.5, respectively.

Clouds typically appear warmer (or brighter) than the night sky at 8–14 µm wavelengths (see Section 2.4). This characteristic allows for a rather simple method of cloud detection, as temperatures warmer than the clear sky background can be identified as clouds. Overall, estimating the background using these findings has not been successful (for an example, see Figure 5.10 in the following section). This is due to the uncertainties associated with estimating A and B. At a given zenith angle θ, it can be shown that sky temperatures predicted using the model (Equation 5.3) have an error ΔT_{sky} that can be calculated from the following:

$$\Delta T_{sky}^2 = \Delta A^2 + \left[\frac{(\Delta \ln \sec \theta)^2}{(\ln \sec \theta)^2} + \frac{\Delta B^2}{B^2} \right] \times (\ln \sec \theta)^2 \times B^2 \,, \tag{5.6}$$

where ΔA is the error in the A parameter (about 0.9 K), ΔB is the error in the B parameter (about 0.4 K), and $\Delta \ln \sec \theta$ is the error in $\ln \sec \theta$ (assumed to be negligible). Unless θ is rather large (greater than 85°), the dominant term in the expression is due to the uncertainty in A. Whilst the value of A determines the offset of the clear sky function (the assumed sky temperature at the zenith), the value of B instead determines the overall "shape". Moreover, the temperature calibrations for our cameras typically have uncertainties of about 2 K. Since the calibrations depend on the temperature of a camera's sensor, which remains fairly constant throughout a given scan, the temperature readings for each of the pixels are affected in the same way. Consequently, only the A fit parameter is influenced by the accuracy of the temperature calibration, as the overall "shape" of the temperature profile remains unchanged.

Another consideration is that because our models have been developed using data collected on clear nights (comparatively drier atmospheres), it may be possible that they cannot properly account for moist atmospheres, which are common when clouds are present. Although estimating clear sky temperatures using these results has not been successful, the models contain some information that can be useful for identifying clouds (see Section 5.2).

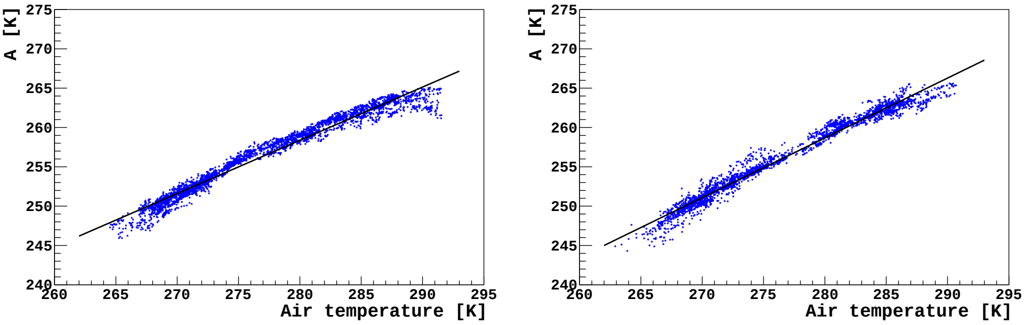

(a) Los Leones **(b)** Los Morados

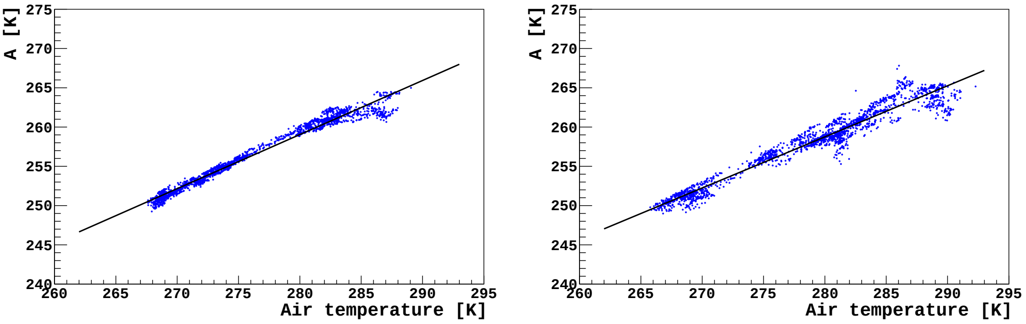

(c) Loma Amarilla **(d)** Coihueco

Figure 5.7: Fit parameter A (in Equation 5.3) plotted against the air temperature. Results are for each of the Auger cloud cameras, and have been obtained using data from multiple scans. Air temperatures can be determined using our cameras, and taken as the sky temperature near the horizon (where the atmosphere is thickest, and resembles a black body). The relationships have been parameterised using linear functions (see Equation 5.4). Fits to the data are indicated by the black lines, and the fit parameters listed in Table 5.1.

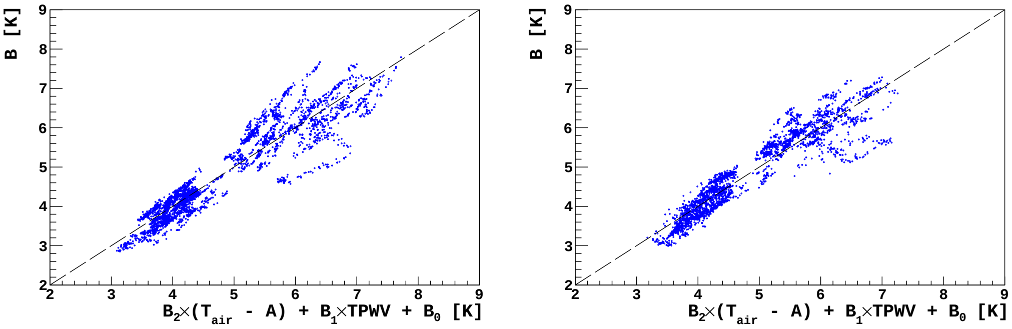

(a) Los Leones **(b)** Los Morados

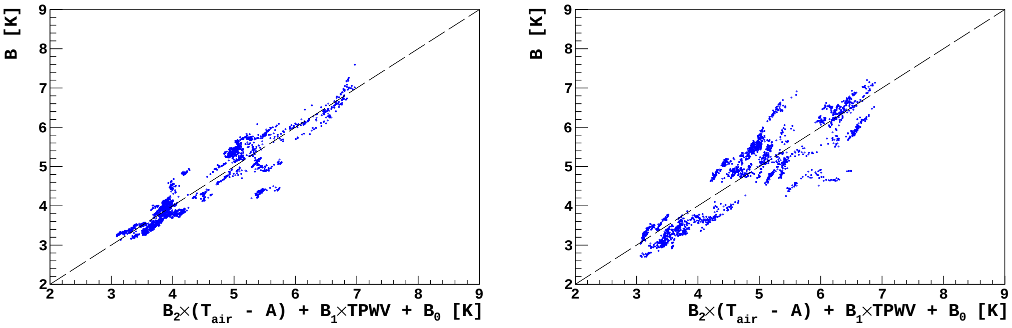

(c) Loma Amarilla **(d)** Coihueco

Figure 5.8: Fit parameter B (in Equation 5.3) plotted against predicted values of B for each of the Auger cloud cameras. Predictions are made using Equation 5.5 and the parameters listed in Table 5.2. Quantities used to estimate B are the air temperature T_{air} (the sky temperature at the horizon in Kelvin), the fit parameter A (in Equation 5.3 also in Kelvin), and the total precipitable water vapour (TPWV), which is the depth of water vapour in an atmospheric column, if all the water precipitated as rain (expressed here in millimetres). One-to-one relationships are indicated by the dashed black lines.

5.2 Detecting clouds with our infrared cameras

The focus of the previous section had been to predict clear night sky temperatures within the band-pass of our cameras. Unlike the vertical atmosphere (which is optically thin on a clear night), the majority of clouds resemble black bodies (see Section 2.4) in the long-wave infrared. Consequently, most clouds have (warm) effective temperatures that stand out against the (cool) clear sky background. This characteristic helps us to identify clouds in our thermal images. However, the process is not necessarily straightforward as water vapour and carbon dioxide are also strong emitters (and absorbers) at these wavelengths. As a result of the increasing optical depths of these particular gases at larger zenith angles, sky temperatures appear warmer towards the horizon (see Equation 5.3), which makes it more difficult to distinguish a cloud's thermal emission from that of a clear sky. This can be problematic, as the main objective of our infrared cameras is to monitor for clouds within the Auger fluorescence detectors' fields of view, which approximately cover between 60°–88° zenith angles. In this section, I outline the procedure that I use to identify clouds in our thermal images.

5.2.1 Method

When analysing the cloud camera images, there are several steps that I use to process the data. These have been outlined in Section 4.6. Figure 5.9 shows an example of a panorama captured with the Los Leones cloud camera after the data has been processed. Clouds can be identified visually as they stand out against the clear sky background. The main goal of this section, however, is to automate their detection.

A rather simple method of cloud detection is to subtract the clear sky background from the data. Given that clouds appear warmer than the background, the residual signal can often be attributed to clouds. It is more appropriate, however, to set some threshold above the background in order to reduce the risk of falsely identifying clouds. As an example, the need for a threshold can be seen in Figure 5.3a (of the previous section) due to the spread in pixel values at any given zenith angle (about 1.5 K). I usually set a temperature threshold 3.5 K warmer than the background to detect clouds, although the threshold is lowered when possible in order to increase the sensitivity. What ultimately determines the threshold is the spread in pixel values that correspond to the clear sky (often about 1.5 K), however, a threshold slightly warmer acts as a "buffer" to avoid false cloud detections.

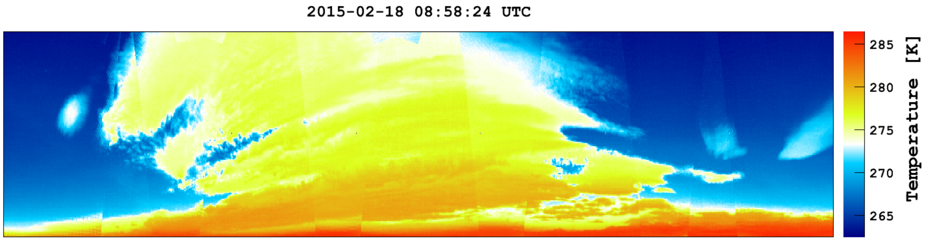

Figure 5.9: A panorama captured during a full-sky scan performed with the Los Leones infrared camera at night-time. Clouds stand out in the images because they have warmer effective temperatures than the clear sky. The panorama has been cropped to show the approximate field of view of the fluorescence detector.

In Section 5.1, an equation for predicting clear sky temperatures T_{sky} (the background) at different zenith angles θ had been expressed as [53]:

$$T_{sky}(\theta) = A + B \times \ln(\sec\theta) \ , \qquad (5.3)$$

where empirical formulas for calculating A and B had been given by Equation 5.4 and Equation 5.5, respectively.

Another way of representing the data in Figure 5.9 is to plot the sky temperatures against zenith angle (see Figure 5.10). The warm bumps in the temperature profile are associated with clouds, which are easily distinguished from the smoother profiles corresponding to clear skies (for example, see Figure 5.3a). In addition, the predicted temperature profile of a clear sky based on the results of the previous section is shown by the red curve. Unless the sky is completely overcast at a given zenith angle, the coolest sky temperatures correspond to a clear sky. Consequently, it can be seen that this particular clear sky profile poorly describes the data. Whilst the B parameter (in Equation 5.3) tends to determine the "shape" of the profile, the A parameter (also in Equation 5.3) instead controls its offset. A major problem with predicting clear sky temperatures using the model, in general, has been the estimations of A (see Section 5.1.1). This is because the residual signal is most affected by this parameter.

In order to avoid incorrectly estimating clear sky temperatures, I have found that it is often better to attempt to locate clear sky regions directly from the data instead. This usually corresponds to the lower bound of pixel values (for example [101]) such as in Figure 5.10. A simple method that can be used to determine the lower bound computationally is to first group the pixel values by zenith angle, and then identify the minimum value within each

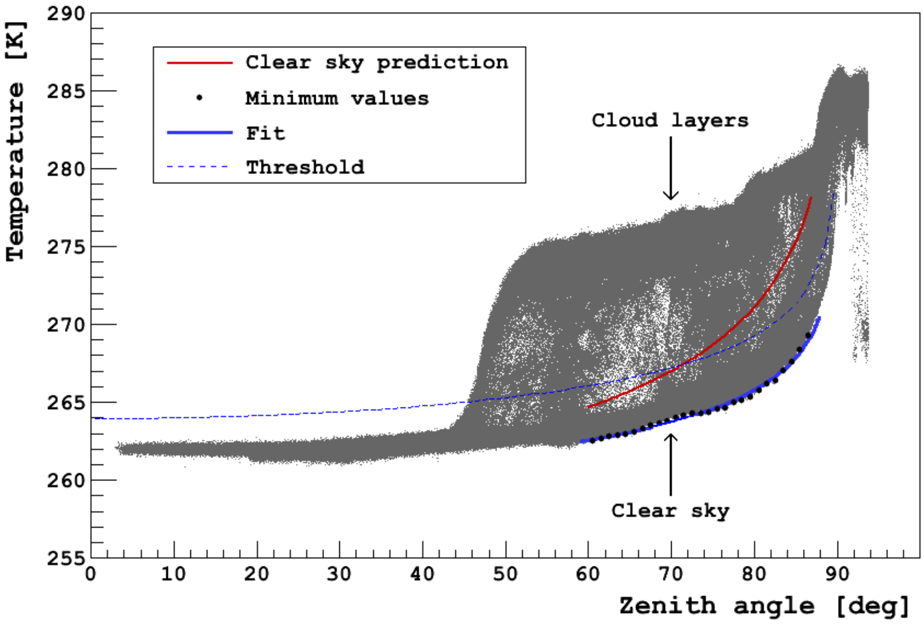

Figure 5.10: Measurements made with the Los Leones infrared camera (the same full-sky scan as Figure 5.9). A temperature reading for each pixel has been plotted against zenith angle. The lower bound of pixels corresponds to the clear night sky background, whereas the warm bumps in the profile are associated with clouds. Predicted clear sky temperatures (based on the results in Section 5.1) as a function of zenith angle are indicated by the red curve, but do not describe the data well. Minimum pixel values between 60°–87° zenith angles are indicated by the black circle markers, and correspond to regions of clear sky. A fit to the minimum values using Equation 5.3 is shown by the blue solid curve. The blue dashed curve indicates a temperature threshold (3.5 K warmer than the background, and extrapolated to span 0°–89.5° zenith angles) that can be used to identify clouds. Pixels viewing clouds generally lie above the threshold.

bin. One of the challenges with using the minimum values, however, is that the results can be rather sensitive to occasional outliers in the data (pixels with much cooler temperatures than expected). The procedure that I use to remove the outliers is to initially bin the data at 0.2° zenith angle intervals. The data are then grouped into larger 1° bins, with each set containing the 5 minimum values. An average value for each set is calculated, and provided that the standard deviation of those values is small (less than 1 K), it will be used in the subsequent analysis. Taken together, for each degree of zenith angle, a single value represents the minimum sky temperature which usually corresponds to the clear sky. Since we are particularly interested in identifying clouds within the Auger fluorescence detectors' fields of view, only minimum values for zenith angles between 60°–87° are used. This is also where Equation 5.3 best fits our cloud camera data (see Figure 5.3 in the previous section). The minimum values for the data in Figure 5.10 are indicated by the black circle markers.

In most cases, the background can be estimated from the minimum values. This is because there is usually some region of clear sky seen by our cameras at any given zenith angle. In fact, when our cameras perform a full-sky scan (spanning 360° azimuth angles), there is an even greater chance of locating clear skies. An example of fitting Equation 5.3 to the minimum values in Figure 5.10 in order to estimate the background is shown by the blue solid line. A temperature threshold 3.5 K warmer than the background is indicated by the blue dashed line, and can be used to identify most clouds. Pixels with temperature readings above the threshold can be identified as viewing cloud, whereas pixels with values below the threshold are instead identified as viewing clear sky. However, it can be seen that some of the pixels viewing cloud are below the chosen threshold (see Figure 5.10 on the previous page). The resultant information is illustrated in Figure 5.11, which shows cloud (light grey) and clear sky (black) data for the same panorama. It can be seen that the clouds identified using this algorithm agree well with those seen through visual inspection of the original data (compare with Figure 5.9). Accurately identifying clouds near the horizon (elevations less than about 3°) is difficult using this method. Such clouds at low elevations are also often hard to see in the original thermal images.

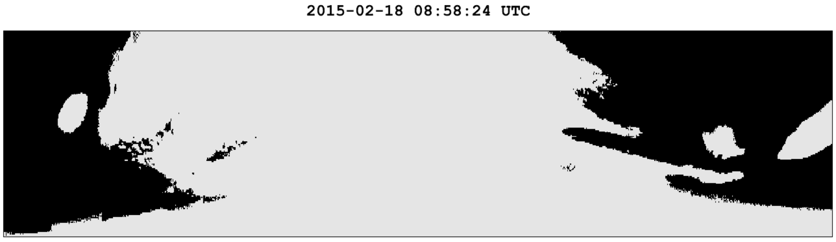

Figure 5.11: An example of cloud layers (light grey) and regions of clear sky (black) identified in the cloud camera data (the same panorama as in Figure 5.9). Camera pixels with temperature readings 3.5 K warmer (the threshold) than the clear night sky (the background) have been identified as viewing cloud.

Although the lower bound often corresponds to clear sky temperatures, it is not appropriate to use the minimum values when the sky is overcast. Figure 5.12 shows another example of sky temperatures plotted against zenith angle using data from a different cloud camera scan. Similarly, the black circle markers indicate the minimum values between 60°–87°. However, the warm values between 67.5°–86.5° are due to clouds rather than clear sky. Consequently, the background signal cannot be estimated using all of the data points.

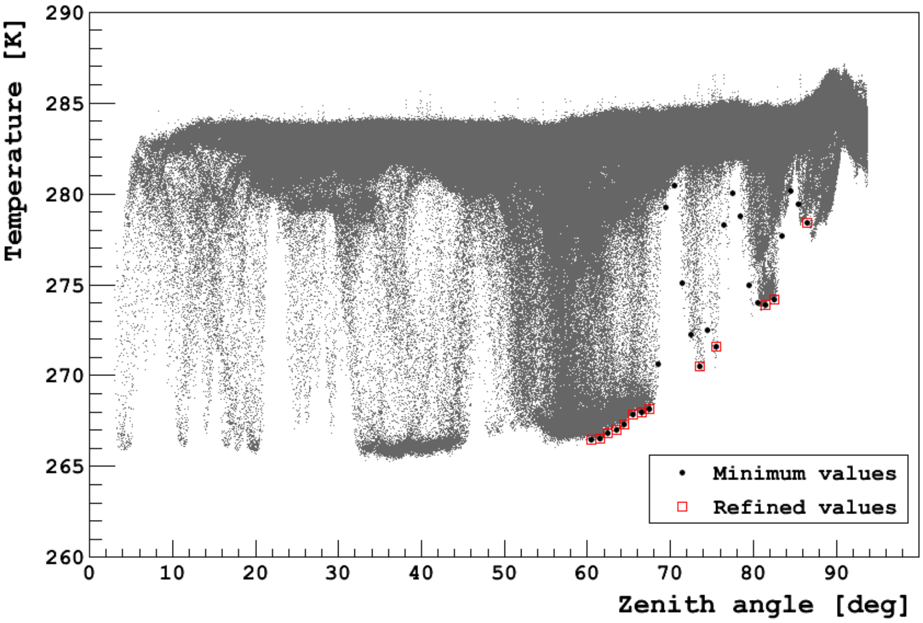

Figure 5.12: Measurements made with the Los Leones infrared camera. Data have been collected from a full-sky scan that began at 2015-01-11 07:52:31 UTC, and consists of 19 individual images. A temperature reading for each pixel has been plotted against zenith angle. The lower bound of pixels corresponds to the temperature of the clear night sky background. The upper bound instead corresponds to the cloud base temperatures. Minimum values for each degree of zenith angle between 60°–87° are indicated by the black circle markers. There are several warm regions between 67.5°–86.5° due to clouds rather than clear sky. Those values can be rejected from the analysis using an iterative method. The remaining values that have been refined using the method are indicated by the red square markers. The background can instead be estimated by fitting a function (Equation 5.3) to the refined values.

I use an iterative technique to refine the data. The main objective is to reject data points corresponding to clouds, whilst maintaining data associated with the clear sky. The first step used to test whether data may be affected by clouds is to cycle through the minimum values in order of increasing zenith angle. The change in sky temperature with zenith angle can be calculated between each point (the numerical derivative). When a data point corresponding to a clear sky (a clear point) is followed by one that corresponds to cloud (a cloudy point), the derivative will be positive due to clouds appearing warmer than clear skies. On the other hand, the derivative will be negative when a cloudy point is followed by a clear point. Given that clear sky temperatures also vary with zenith angle, the data cannot be refined based on this criteria alone.

The approximate change in clear sky temperature T_{sky} with zenith angle θ (in radians) can be calculated by differentiating Equation 5.3. It can be shown that the analytical derivative is given by:

$$\frac{\mathrm{d}T_{sky}}{\mathrm{d}\theta} = B \times \tan\theta\,, \tag{5.7}$$

where B can be estimated from Equation 5.5 and the parameters listed in Table 5.2. Data points are retained in this analysis provided that the numerical derivative (calculated using neighbouring points) is somewhat similar to the analytical value (calculated using Equation 5.7). I typically require the absolute difference between both quantities to be less than 0.4 K (the approximate uncertainty in B, see Figure 5.6b). Since it has been difficult to find a single value that can be used in all situations, I adjust the parameter when necessary.

A summary of the steps used to refine the data are listed below:

1. Loop over the minimum values in order of increasing zenith angle.

2. Calculate the change in sky temperature with zenith angle numerically between each data point.

3. Calculate the change in sky temperature with zenith angle analytically using Equation 5.7. The equation is evaluated for a zenith angle halfway between both data points.

4. Determine the absolute difference between the numerical and analytical (from Equation 5.7) quantities. There are several outcomes depending on the result.

    - Keep both points provided that the difference is small (typically less than 0.4 K). This suggests that both values correspond to regions of clear sky.

- When the difference is large and the numerical derivative is positive, it suggests that the second point (at a greater zenith angle) corresponds to a cloudy region of the sky. Remove the second point from the analysis.
- When the difference is large and the numerical derivative is negative, it suggests that the first point (at a lower zenith angle) corresponds to a cloudy region of the sky. Remove the first point from the analysis.

5. Repeat the process until each point satisfies the criterion for a clear sky (i.e. the numerical derivative and analytical derivative are in good agreement).

Figure 5.12 shows an example of applying this iterative technique to the cloud camera data. The minimum values that have been refined using the technique are displayed by the red square markers. Following the same procedure as before, the background can be estimated by fitting Equation 5.3 to the remaining minimum (refined) values. Clouds can subsequently be identified by setting a temperature threshold above the background.

In the previous example, the minimum clear sky values remained after applying the iterative technique to the data. On the other hand, fewer values would remain when the sky is completely overcast. This is because overcast sky temperatures do not vary appreciably with zenith angle (see the upper bound in Figure 5.12). As a result, many data points will be removed from the analysis. However, since clear sky temperatures at lower zenith angles (less than $\sim 70^\circ$) also remain fairly constant, not all of the data points can be removed using this technique. To help automate the process, overcast skies are identified when there are fewer than 12 minimum values remaining (i.e. fewer than $\sim 44\,\%$ of the minimum values are believed to correspond to regions of clear sky). In addition, small values of B (which relate to a small temperature difference between 60° and 87° zenith angles) obtained from fitting Equation 5.3 to the lower bound can also help to identify overcast skies. On clear nights, values of B were found to be greater than $2\,\mathrm{K}$ (see Figure 5.8). A value of B less than $2\,\mathrm{K}$ can instead be used to infer an overcast sky. Both criteria are often used when analysing the cloud camera data, although some manual effort is still required to check the performance. An overview of the procedure concludes this section.

5.2.2 Concluding remarks

In this section, a method for detecting clouds in our thermal images was developed. Since clouds typically appear warmer than clear night skies at infrared wavelengths, they usually stand out in our cloud camera data. A simple method that can be used to identify clouds is to subtract the background (clear sky) signal from the data. The residual signal can often be attributed to clouds. In practice, I have found that it is better to determine the background from the measurements rather than estimating it from the empirical models found in Section 5.1.

The scanning routines performed by our cameras cover a wide field of view. As a result, they often observe regions of clear sky even on nights with partial cloud cover (unless the sky is completely overcast). Consequently, when plotting sky temperatures as a function of zenith angle, the clear sky background corresponds to the lower bound of pixel values (see Figure 5.10). I fit a clear sky model (Equation 5.3) to the lower bound between 60°–87° zenith angles. Not only does the model best fit the data in that range, but because it is also the approximate field of view of the Auger fluorescence detectors, it is where we are most interested in identifying clouds.

When the sky is overcast at a given zenith angle, the coolest temperature reading instead corresponds to cloud rather than the clear sky. In order to estimate the clear sky background, the data can be refined using an iterative technique to reject measurements affected by clouds. Despite this, simply fitting to the lower bound works well in most cases. Taken together, I have a method that can help to refine the data when required.

A temperature threshold warmer than the clear sky background can be used to identify clouds. Pixels with temperature readings warmer than the threshold can be identified as viewing cloud, whereas pixels with values cooler than the threshold can be identified as viewing clear sky. I generally set the temperature threshold to be 3.5 K warmer than the background, although the threshold is lowered when possible in order to increase the sensitivity to cloud detection. Occasionally, the threshold must be raised to avoid false detections of cloud. The threshold is ultimately determined by the spread in the pixel values corresponding to the clear sky background. Although there is still some manual effort involved in adjusting the threshold, the values are typically only adjusted by about ±1 K. It is important to note that without the temperature calibration, finding an appropriate threshold is more laborious as the pixel outputs (instead in ADC counts) also depend on the temperature of the camera's sensor, which varies throughout a given night. Examples that demonstrate why we wish to compensate for the temperature of the sensor may be found in Section 4.5.

Our thermal imaging cameras are most sensitive to low, thick clouds (see Section 2.4). Not only are clouds at lower altitudes at warmer absolute temperatures (emitting a greater radiant flux), but when they are optically thick they resemble a black body. As a result, they often stand out in our images. On the other hand, clouds at higher altitudes are more difficult to detect because they are cooler, and typically have lower emissivities. In addition, depending on the humidity, they can be obscured by the intervening atmosphere. In fact, it had been shown in Section 2.4 that cirrus clouds have emission spectra that closely resemble that of a clear sky with water vapour.

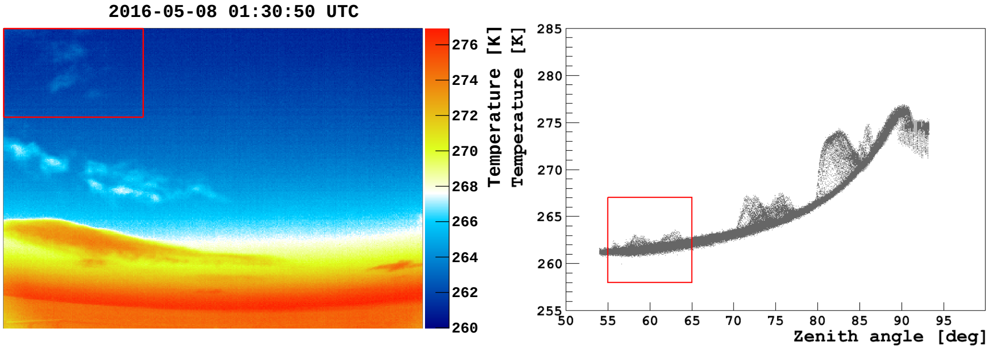

(a) **(b)**

Figure 5.13: **(a)** An image captured with the Loma Amarilla infrared camera when clouds were present. One of the cloud layers (within the red rectangle) appears slightly warmer (about 1 K) than the clear sky background, and can only be faintly seen. **(b)** The temperature reading for each pixel plotted against zenith angle using the same data. The clear night sky background is the coolest region of the sky, and can be seen by the lower bound of pixel values. The same cloud layer highlighted in (a) is contained within the red rectangle, and lies just above the background.

To demonstrate the limitations of cloud detection using our infrared cameras (and the current analysis), consider the data presented in Figure 5.13. Figure 5.13a shows an image of several cloud layers captured with the Loma Amarilla cloud camera. The same clouds can also be seen in Figure 5.13b as warm bumps in the temperature profile. Interestingly, the cloud layer between 55°–65° zenith angles (within the red rectangles) appears to be slightly warmer (about 1 K) than the clear night sky background, and can only be faintly seen through visual inspection of the original image (Figure 5.13a). Consequently, a low temperature threshold would be required to detect this particular cloud. In practice, however, there are difficulties associated with using such low thresholds because there is a greater chance of falsely identifying clouds. In this case, the cloud layer's infrared flux contribution to the overall sky brightness (which relates to the effective sky temperature) is rather low, which suggests that the cloud is optically thin and/or at high altitudes. Another example of this may be seen in Figure 5.14, which shows an image captured with the Los Leones cloud camera. It can be seen in Figure 5.14a that a cloud layer faintly covers the camera's field of view. However, the layer of cloud is not obvious when inspecting the temperature profile (see Figure 5.14b). This is because the variation in effective sky temperature (when observing the cloud) with zenith angle is similar to that of a clear night sky (i.e. approximately $\propto \ln(\sec\theta)$). In practice, it has been easier to identify thin clouds such as Figure 5.14 through visual inspection rather than using the current analysis. However, as noted in [46], optically thin clouds do not always have a noticeable effect on individual extensive air shower events. Thin clouds can instead be detected using the CLF and XLF (see Section 3.5) at the Pierre Auger Observatory.

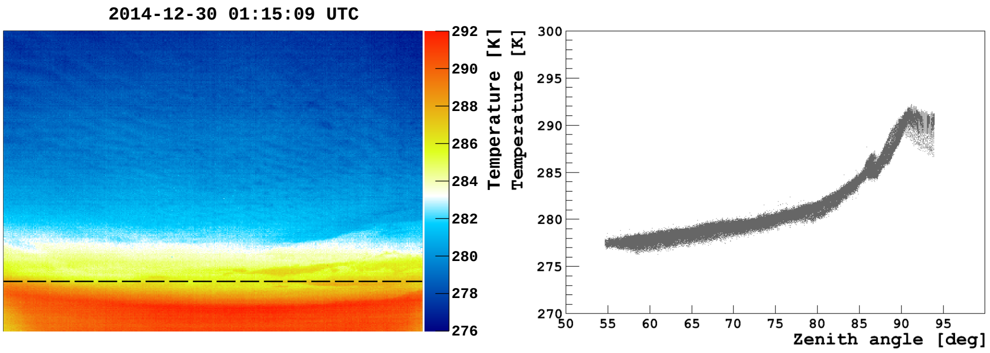

(a) **(b)**

Figure 5.14: **(a)** An image captured with the Los Leones infrared camera. A layer of cloud at a cool effective temperature can faintly be seen in the camera's field of view. The Auger fluorescence detectors' have fields of view that approximately cover between 60°–88° zenith angles. A black dashed line at about 88° is shown, whereas 60° is located near the top of the image. **(b)** The temperature reading for each pixel plotted against zenith angle (θ) using the same data. Note that the variation in effective sky temperature (when observing cloud) with zenith angle is similar to that of a clear night sky (i.e. approximately $\propto \ln(\sec\theta)$).

Detecting clouds near the horizon is also difficult using this method. One of the main reasons is due to the increased opacity of the atmosphere at larger zenith angles. Another problem is that the function used to predict clear sky temperatures (Equation 5.3) does not describe our cloud camera data well at larger zenith angles (larger than $\sim 87°$). This could perhaps be improved by simulating the clear night sky background with a radiative transfer model (such as MODTRAN) using atmospheric profiles rather than relying on an empirical formula. In addition, it may also be possible to lower the cloud detection threshold, which would also increase our sensitivity to optically thin and high altitude clouds (such as those in Figure 5.13 and Figure 5.14). This would also help to reduce the manual effort involved in adjusting the thresholds. Some preliminary results using this approach are the focus of Section 5.4. With the current analysis, the best results are generally obtained at elevations greater than about 3°.

A general outline of the current analysis is listed below:

1. Process the cloud camera data (i.e. apply flat-field corrections to the images, and convert pixel outputs into temperature readings). The details may be found in Section 4.6.
2. For each degree of zenith angle between 60°–87°, determine the minimum pixel values. The coolest temperature readings usually correspond to the clear night sky background. Outliers are removed during this process.
3. When the lower bound does not correspond to the background, measurements influenced by clouds can be removed from the analysis using an iterative technique. Fewer than 12 remaining points can help identify overcast skies.
4. Fit a function (Equation 5.3) to the lower bound that describes the change in sky temperature with zenith angle. A small value for the B fit parameter (less than $2\,\mathrm{K}$) can identify overcast skies.
5. Set a temperature threshold about $3.5\,\mathrm{K}$ warmer than the clear sky background. Pixels with values warmer than the threshold are identified as viewing cloud, whereas pixels with values cooler than the threshold are identified as viewing clear sky. The threshold is adjusted (typically by about $\pm 1\,\mathrm{K}$) to increase the sensitivity to cloud detection, whilst minimising false detections of clouds.

Given that there is still some manual effort required, I analyse the cloud camera data only when the Auger fluorescence detectors are operating.

This method of cloud detection is used to generate cloud "masks" (see Section 5.3), which provide the fraction of cloud contained within each of the Auger fluorescence detector pixels. The cloud masks after January 2015 have been produced using this technique. Due to experimentation with the camera settings (see Appendix B), data recorded between 2013 (when the cameras were installed) through to March 2014 had to be analysed separately. A somewhat similar method of cloud detection was performed by fitting a function to the lower bound of pixel values (in ADC counts). However, it was not straightforward to determine flat-field corrections and temperature calibrations during this period. Cloud masks between March 2014 through to January 2015 had been produced using a separate analysis performed by Trent Grubb, a current PhD. student at the University of Adelaide.

5.3 Cloud masks

Clouds can affect measurements made by the Auger fluorescence detectors (FDs). This is because both nitrogen fluorescence and Cherenkov light from extensive air showers can be scattered by clouds (see Section 1.2.3). Correcting for these effects, however, is not a straightforward task. One of the major challenges is due to the uncertainty in a cloud's scattering properties [24], which depends on its composition, and can be difficult to obtain. Since it is difficult to account for the influence that clouds can have on measurements made using the fluorescence technique, potentially affected data are often removed from analyses.

There are various instruments used for cloud detection at the Pierre Auger Observatory (see Chapter 3). Most relevant to my research have been the thermal imaging cameras, which are dedicated to monitoring for clouds located within the FD fields of view. Our cameras are installed at the Los Leones (LL), Los Morados (LM), Loma Amarilla (LA), and Coihueco (Co) fluorescence detector sites (see map, Figure 3.1). The main objective of our infrared cloud cameras is to generate ”masks” that provide the cloud cover contained within the fluorescence detectors' pixels. This information is stored into the cloud camera database, which is one of the atmospheric databases maintained at the Observatory. In this section, I describe the details of the cloud masks.

5.3.1 Format

The Auger fluorescence telescopes are designed to focus atmospheric fluorescence light onto a specialised camera (see Section 1.2.2). Each fluorescence camera consists of 440 hexagonal photomultiplier tubes (the FD pixels), which are arranged into 22 rows and 20 columns. There are six telescopes located at each of the main fluorescence detector sites, and the individual FD pixels have an approximate 1.5° field of view. Figure 5.15 shows a panorama taken at night-time with the Los Leones infrared camera when clouds were present. The clouds stand out in the images because they have effective temperatures that are warmer than the surrounding clear sky. In addition, the pointing directions of the 2640 FD pixels for the site are also shown (grey hexagons). A cloud mask can be created by determining the fraction of cloud (cloud cover) contained within each FD pixel.

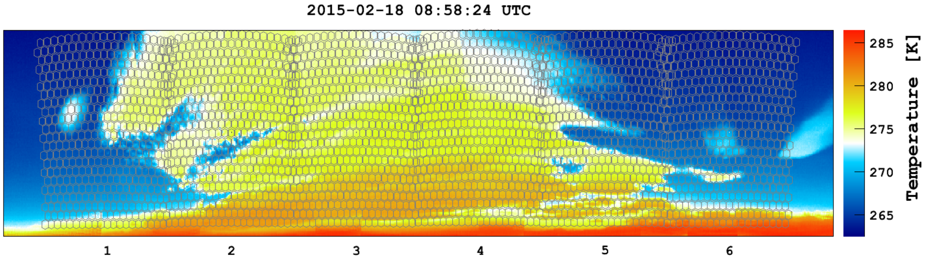

Figure 5.15: A panorama captured with the Los Leones infrared camera at night-time (the same as Figure 5.9). Because clouds typically have effective temperatures warmer than that of a clear sky, they usually stand out in our thermal images. The pointing directions of the Auger fluorescence detector (FD) pixels for the site are also shown (grey hexagons). The fraction of cloud (cloud cover) contained within each FD pixel can be calculated, and used to produce a cloud mask (see Figure 5.16).

Clouds are identified using the method described in Section 5.2. In order to determine the cloud cover within a FD pixel, it is more straightforward to consider circular pixels rather than hexagons. Each FD pixel is defined to have a 0.75° angular radius for this analysis, and contains about 70 cloud camera pixels. For each fluorescence detector pixel, the fraction of infrared camera pixels containing cloud can then be determined, and an index is assigned. There are six indices that are used to represent the amount of cloud within a FD pixel's field of view, which range from 0 (with 0–10 % cloud cover) through to 5 (with 90–100 % cloud cover). The cloud indices that are allocated are listed in Table 5.3, which also includes the corresponding values when data are read from the cloud camera database. The majority of FD pixels are assigned either a 0 or 5 (i.e. they are mostly clear or overcast). Only the pixels containing a cloud's boundaries will contain mixed fractions (for example, see Figure 5.16). Figure 5.16 shows an example of a cloud mask (using the same data as Figure 5.15) that has been created after assigning an index to each FD pixel. This information is written into a database file before it is uploaded into the cloud camera database.

The database filenames are saved using the format `SSYYYYMMDD.cpd`. The first two characters `SS` identify the site (i.e. LL, LM, LA or Co), and are followed by the year `YYYY`, the month `MM`, and the day `DD` the fluorescence detector was operating (an FD shift, where the date corresponds to that of the noon prior to the shift). As an example, the cloud mask information in

Fraction of cloud [%]	Cloud index	Database output
0–10	0	0
10–30	1	0.2
30–50	2	0.4
50–70	3	0.6
70–90	4	0.8
90–100	5	1

Table 5.3: A list of cloud indices that are assigned to the Auger fluorescence detector (FD) pixels. Each index represents the fraction of cloud contained within a FD pixel's field of view. The corresponding values when data are read from the cloud camera database are also shown. Table modified from [22].

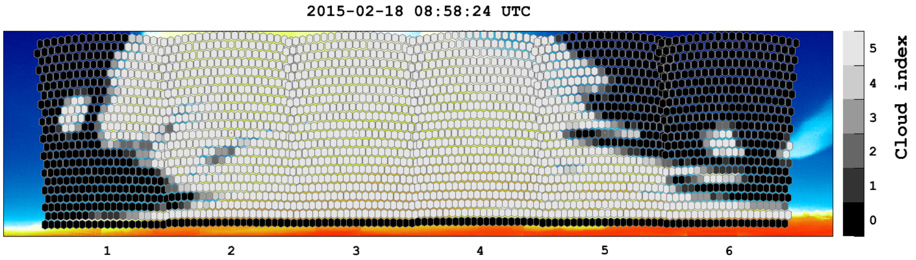

Figure 5.16: A panorama captured with the Los Leones infrared camera (the same as in Figure 5.15). The pointing directions of the Auger fluorescence detector (FD) pixels for the site are also shown (hexagons). Each FD pixel has been assigned a cloud index (see Table 5.3) that represents the fraction of cloud contained within its field of view. Lighter colours (a higher cloud index) correspond to greater amounts of cloud cover.

Figure 5.16 is saved in the file `LL20150218.cpd` along with any other masks for that particular FD shift.

There are multiple lines in each database file. Each line contains the time the camera began its scan (in GPS time), a site identification number (1 for LL, 2 for LM, 3 for LA, and 4 for Co), the telescope identification number (1 through to 6), and are followed by the 440 cloud indices corresponding to that telescope's FD pixels. The ordering of the pixels follows the convention

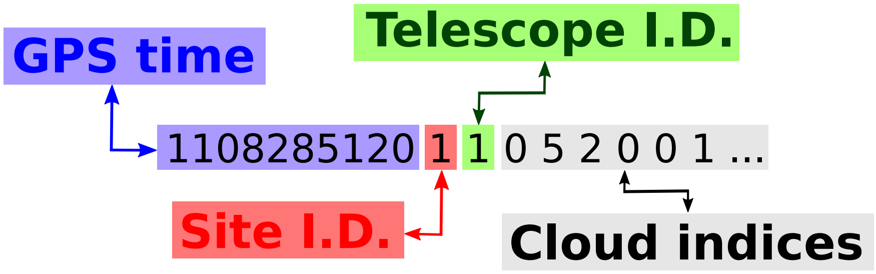

Figure 5.17: An example of the information written into a cloud mask database file. The first entry contains the time the camera began its scan (in GPS time), the second entry identifies the fluorescence detector site (1 to 4), and the third entry identifies the telescope (1 to 6). This is followed by 440 cloud indices corresponding to the cloud cover in the individual fluorescence detector pixels.

used for the Observatory. Figure 5.17 shows an example of the structure for a single line. The next 5 lines would contain cloud mask information for the remaining 5 telescopes (and have the same GPS time). This is followed by another set (6 lines) of cloud mask data for the next time sequence (usually 5 minutes later).

Each set of 2640 cloud indices are also assigned validity times (details discussed in [22]). The times consist of 2 numbers, which act as lower and upper bounds (in GPS time) that a particular index is valid. Since our cloud cameras are programmed to record data at 5 minute intervals, the validity times typically extend from 2.5 minutes before to 2.5 minutes after each scan. In the event that data are not recorded (possibly due to a temporary camera malfunction), the validity time may be extended by up to 10 minutes.

In order to minimise the amount of storage space required, the validity times can also be used to compress our data [22]. When the cloud index assigned to a particular FD pixel is the same for neighbouring validity times, the validity period can be extended. This means that only a single entry (with a combined validity period) is needed rather than two identical entries (with two differing validity periods).

The cloud mask data are transferred to the Auger Computer Centre at Wuppertal. The information can be accessed by other members of the collaboration from the cloud camera database, which is one of the Auger $\overline{\text{Off}}\underline{\text{line}}$ atmospheric databases (database outputs listed in Table 5.3). Another way is through the Pierre Auger Observatory FD Cloud Mask Viewer [1], which has

---

[1] The Pierre Auger Observatory FD Cloud Mask Viewer can be accessed from `http://www.physics.adelaide.edu.au/astrophysics/Auger/CloudCams/`.

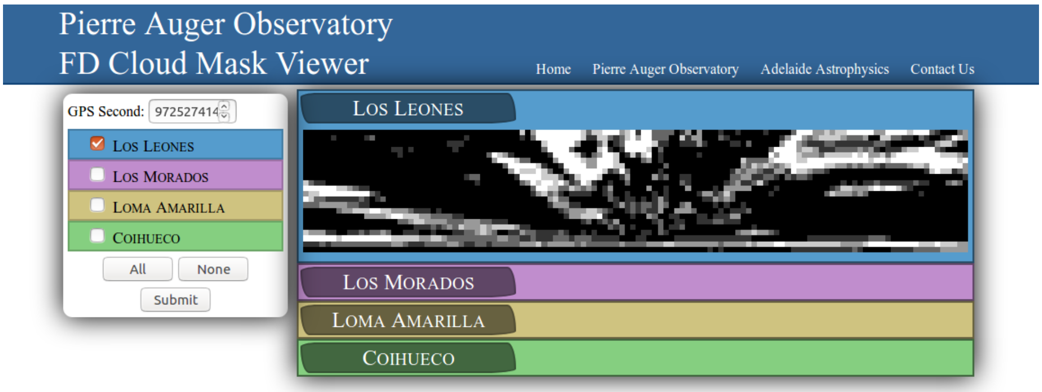
GPS Second: 972527414

☑ Los Leones

Figure 5.18: The Pierre Auger Observatory FD Cloud Mask Viewer may be accessed from `http://www.physics.adelaide.edu.au/astrophysics/Auger/CloudCams/`. The website shows the cloud cover contained within the individual fluorescence detector (FD) pixels, which it accesses from the cloud camera database. Lighter colours correspond to greater amounts of cloud cover. Users can specify the GPS time and what FD site(s) they are interested in. Created by Mathew Cooper and Benjamin Whelan.

been created by Mathew Cooper and Benjamin Whelan. The website allows users to specify the GPS time and select what FD sites they are interested in. Data is read from the cloud camera database, and the corresponding cloud masks are displayed (see Figure 5.18).

5.4 Simulating thermal radiation from clear night skies

The majority of clouds have (warm) effective temperatures that stand out against the (cool) clear night sky background at infrared wavelengths. Knowledge of the background can be used to distinguish a cloud's thermal emission from that of a clear sky. Currently, the background signal is estimated by fitting an empirical function (Equation 5.3) to clear sky regions, which are identified directly from the camera data (see Section 5.2). Rather than relying on an empirical formula, another approach would be to simulate the clear sky background (for example [102][103]) using knowledge of the current atmospheric conditions. In this section, I outline some preliminary results using this technique.

Thermal emission from a clear night sky can be simulated using a radiative transfer model such as MODTRAN. Some details regarding the program may be found in Appendix A. One of the features in MODTRAN is that it allows users to define an atmosphere using height-dependent atmospheric profiles. This information can, for example, be obtained using data collected from a radiosonde. What is particularly useful for our research is that we have access to temperature, pressure, and humidity profiles from the Global Data Assimilation System (GDAS, see Section 3.6) that are available at 3-hour intervals. An example of using GDAS data as inputs for MODTRAN may be found in Figure A.2 of Appendix A.

We are specifically interested in thermal radiation within the band-pass of our detectors. Each cloud camera is housed within a weather shield and views the sky through a protective window. Consequently, infrared flux from the atmosphere as well as the window material both contribute to the overall signal. The spectral response for the cameras and the band-pass specification of the window have been shown in Figure 3.3 of Chapter 3. The band-pass for each system is a little ambiguous, and the best results were found by setting the camera transmittance to zero above 17 µm and the window transmittance to zero above 14.6 µm. The spectral radiance (in $\mathrm{W\,sr^{-1}\,m^{-2}\,\mu m^{-1}}$) a camera receives R_λ is the sum of the atmospheric and window flux contributions. That is:

$$R_\lambda = R_{atmos_\lambda}\tau_{window_\lambda}\tau_{camera_\lambda} + R_{window_\lambda}\tau_{camera_\lambda}\,, \tag{5.8}$$

where for a given wavelength λ, R_{atmos_λ} is the atmospheric radiance, R_{window_λ} is the window radiance, τ_{window_λ} is the transmittance of the window, and τ_{camera_λ} is the transmittance of the camera.

For a given wavelength λ, the window radiance depends on its emissivity $\epsilon_{window_\lambda}$ and temperature (T_{window} in Kelvin), and can be expressed as:

$$R_{window_\lambda} = \epsilon_{window_\lambda} B_\lambda \left(T_{window}\right)\,, \tag{5.9}$$

where B_λ is the spectral radiance of a black body given by Planck's law. Assuming that the window is at ground temperature (T_{ground} in Kelvin), and since $\tau_{window_\lambda} = 1 - \epsilon_{window_\lambda}$, the spectral radiance the camera receives is given by:

$$\begin{aligned} R_\lambda = {} & R_{atmos_\lambda}\tau_{window_\lambda}\tau_{camera_\lambda} \\ & + \left(1 - \tau_{window_\lambda}\right) B_{window_\lambda}\left(T_{ground}\right)\tau_{camera_\lambda}\,. \end{aligned} \tag{5.10}$$

The total radiance can be calculated by integrating Equation 5.10 across all wavelengths.

Figure 5.19a shows the simulated spectral radiance of a clear sky for a sensor viewing the horizon. The atmosphere has been defined using GDAS temperature, pressure, and humidity data at the Pierre Auger Observatory for 2015-02-11 03:00:00 UTC. In addition, the spectral radiance of a black body at ground temperature (at 292 K) is shown. The infrared flux contributions due to the atmosphere and window material are both indicated. The total radiance can be calculated by integrating Equation 5.10 over all wavelengths, and is found to be $\sim 59.8\,\mathrm{W\,sr^{-1}\,m^{-2}}$. Similarly, Figure 5.19b shows the atmospheric and window flux contributions for a sensor viewing the zenith, which yields a lesser total radiance of $\sim 33.5\,\mathrm{W\,sr^{-1}\,m^{-2}}$.

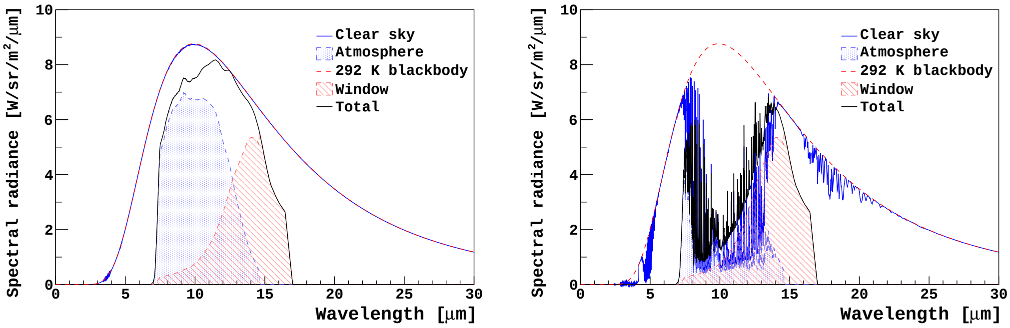

(a)

(b)

Figure 5.19: Atmospheric radiance from a clear sky reaching the ground (blue line). Simulated by the author with MODTRAN using a user-defined atmosphere. The temperature of the Earth's surface (at 1401 m a.s.l) is 292 K in this model, and the radiance of a black body at that temperature is shown (red dashed line). A camera would receive infrared flux contributions from the atmosphere (shaded blue) and the window material (at surface temperature and shaded red). The total spectral radiance within a camera's band-pass is indicated by the black line. **(a)** Simulated for a camera viewing the sky at the horizon. The total flux is $\sim 59.8\,\mathrm{W\,sr^{-1}\,m^{-2}}$. **(b)** Simulated for a camera viewing the sky at the zenith. The total flux is $\sim 33.5\,\mathrm{W\,sr^{-1}\,m^{-2}}$.

In order to convert the camera radiance into an effective sky temperature, a calibration curve can be created. This can be achieved by simulating the radiance within a camera's band-pass for known field of view temperatures. The radiance is calculated by simulating a camera which views the horizon, such as in Figure 5.19a. Since the atmosphere is opaque near the horizon, it resembles a black body at ground temperature. Calibrating the camera using black bodies of known temperatures is similar to what would typically be performed in a laboratory. I have chosen to simulate atmospheric conditions using data from 102 GDAS profiles. The data were selected from a catalogue of clear nights that I had constructed (listed in Table 4.1) in order to make comparisons to radiometer measurements at a later stage. The selection spans a wide range of typical surface temperatures at Auger.

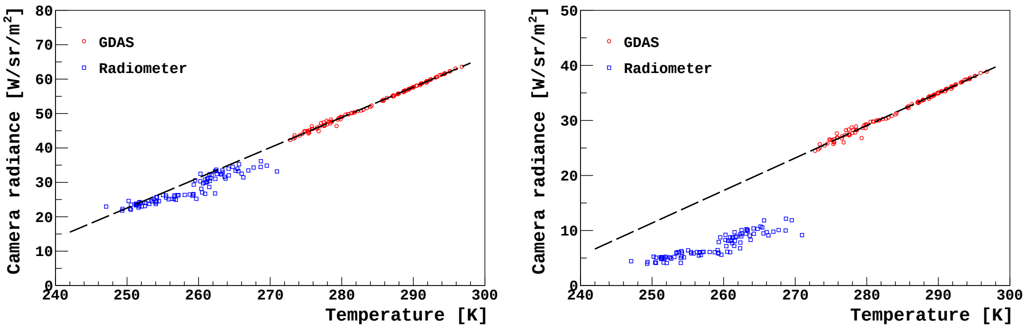

(a) (b)

Figure 5.20: **(a)** A calibration curve showing the simulated radiance within a camera's band-pass plotted for different scene temperatures. Simulated with MODTRAN using a selection of GDAS atmospheric profiles at times when the sky was believed to be clear. The calibration has been performed by simulating the radiance for a camera viewing the horizon (where the atmosphere is opaque and resembles a black body) and comparing the results to GDAS surface temperatures (at 1401 m a.s.l.). The results are indicated by the red circle markers, and a fit to the data is shown by the black dashed line (see Equation 5.11). In addition, coincident radiometer sky temperature measurements are indicated, which can be compared to the radiance simulated for a camera viewing the zenith (blue square markers). Thermal emission due to the window has been included. **(b)** The same as in (a), however, thermal emission due to the window has not been taken into account.

A calibration curve for our cameras is presented in Figure 5.20a. The plot shows the simulated radiance for a camera viewing the horizon for different GDAS surface temperatures. Although the flux is proportional to the fourth power of an object's temperature (Stefan-Boltzmann law), a linear equation is a good approximation over this temperature range. The equation of the line fitted to the data is given by:

$$R_{cam} = 0.878 \times T - 196.9\,, \tag{5.11}$$

where R_{cam} is the simulated radiance (in $\mathrm{W\,sr^{-1}\,m^{-2}}$) within the camera's band-pass and T is the temperature of the black body (in Kelvin).

Also shown in Figure 5.20a are zenith sky temperatures measured with the single-pixel infrared radiometer at the Central Laser Facility (recorded at

5 minute intervals) that are coincident with the 3-hourly GDAS data. The radiometer had been calibrated in a laboratory using an experimental black body, and a linear signal response to temperatures (over a similar temperature range) had also been used as an approximation. What stands out is that the measurements are rather consistent with the radiance simulated for a camera viewing the zenith. This suggests that our cloud camera systems have a similar spectral response to the radiometers at Auger, which was in fact assumed in Section 4.5 when calibrating our cameras. Interestingly, Figure 5.20b shows that not accounting for the window's thermal emission (only including its transmittance) results in poor agreements with the radiometer data. Without the window emission, simulated camera sky temperatures at the zenith are much cooler than what we observe (typically at $\sim 230\,\mathrm{K}$).

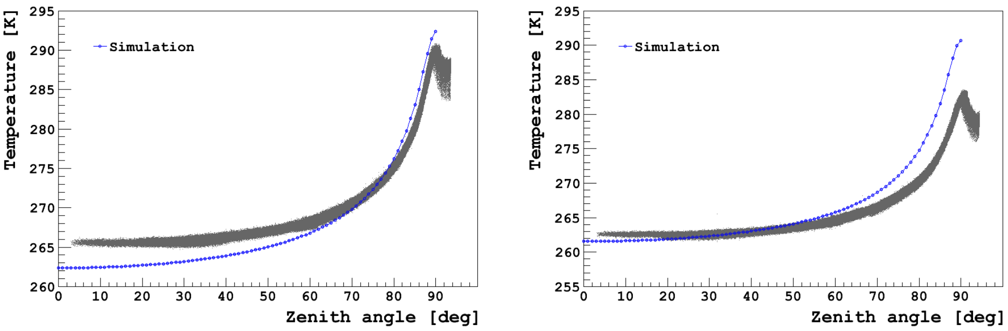

(a) (b)

Figure 5.21: Measurements made by the Los Leones infrared cloud camera on two different clear nights. The temperature reading of each pixel has been plotted against zenith angle θ. For each degree of zenith angle, the radiance within the band-pass of the camera has been simulated using MODTRAN, and converted into an effective sky temperature (shown in blue) using Equation 5.11. The atmosphere has been modelled using height-dependent atmospheric profiles from GDAS. **(a)** A full-sky scan that began at 2015-02-11 01:36:49 UTC. **(b)** A full-sky scan that began at 2016-02-14 06:04:28 UTC.

Using Equation 5.11 allows for a conversion between the simulated radiance and an effective sky temperature. The sky temperatures predicted from the simulation can be compared to real data collected with our cloud cameras. Figure 5.21a shows Los Leones cloud camera measurements from a full-sky scan that began at 2015-02-11 01:36:49 UTC. For each degree of zenith angle, the sky temperature has been computed using MODTRAN loaded with the corresponding GDAS data for that time (the closest data available are at 2015-02-11 03:00:00 UTC). In this case, the simulation agrees reasonably well with the cloud camera observations, especially near the horizon (and only a $\sim 3\,\mathrm{K}$ discrepancy at the zenith). As another example, consider the data presented in Figure 5.21b. The full-sky scan began at 2016-02-14 06:04:28 UTC, and the atmosphere was modelled using GDAS data available at 2016-02-14 06:00:00 UTC. Although the sky temperature simulated at the zenith agrees well with the camera measurements, rather large discrepancies can be seen towards the horizon (about $8\,\mathrm{K}$ at the surface).

There are several reasons that may explain the observed disagreements between the simulations and real data. First, the limited spatial and temporal resolution of the GDAS model can be expected to introduce some uncertainties in the predicted camera radiance and resultant effective sky temperatures. Second, the temperature calibrations developed for our cameras typically have errors of about $2\,\mathrm{K}$. Since the calibrations depend on the temperature of a camera's sensor, which remains fairly constant throughout a given scan, the temperature reading of each pixel will be affected in the same way. This would lead to a systematic offset between the simulated sky temperatures and cloud camera sky temperatures. Lastly, although data collected from local radiosonde launches have been found to be in excellent agreement with the GDAS model at heights above $\sim 5\,\mathrm{km}$ a.s.l. [21], there are some discrepancies that become apparent near the ground. In particular, comparisons to local weather stations in Section 3.6.3 had shown that GDAS tends to overestimate night-time surface temperatures (see Figure 3.16a). Given the potential discrepancies between surface temperatures measured with our cameras and those predicted by the GDAS model, I suspect this to be the main reasons why the simulations do not necessarily match our observations. Although a diurnal variation in the mean accuracy of GDAS ground-level temperatures can be seen (refer to Figure 3.17a), the model minus weather station temperature distributions typically have standard deviations of about $4\,\mathrm{K}$. The rather large spread means that compensating all GDAS profiles for this effect is difficult. Additionally, it is not straightforward to correct the temperature profiles since upper-level measurements agree with radiosonde data.

Because thermal radiation emitted by all atmospheric layers contributes to the infrared brightness seen at the Earth's surface, the temperatures of multiple layers up to about $5\,\mathrm{km}$ a.s.l. (where there are agreements with sonde measurements) likely need to be corrected. However, simply adjusting the GDAS temperature lapse rates below $5\,\mathrm{km}$ a.s.l. such that surface temperatures agree with cloud camera data has not been an adequate correction. Simulated effective sky temperatures for zenith angles above the horizon are instead found to be cooler than what is observed (for example, see Figure 5.22). Increasing the water vapour pressure at different atmospheric layers in the simulations has been studied. Although water vapour can somewhat compensate for the cooler sky temperatures (increased humidity leads to warmer sky temperatures), a clear correction has yet to be found. It is also not straightforward to adjust the GDAS humidity profiles given that water vapour pressure calculations depend strongly on air temperature. Investigating better corrections to the GDAS profiles will be the subject of future work.

Interestingly, this method uses atmospheric profiles in order to simulate the clear night sky background. However, results in the following chapter suggest that GDAS temperature and humidity profiles can be used to infer whether clouds are likely to be present (see Section 6.2). Hence, these profiles do not necessarily reflect completely clear sky backgrounds in which residual signals in our thermal images can be identified as clouds. The MODTRAN atmospheres can be simulated assuming clear sky conditions, even though some of the atmospheric layers may be humid (which may be attributed to clouds). Since a sky with clouds (such as in a cloud camera image) is brighter at infrared wavelengths than a sky without clouds (simulated), pixels viewing cloud would still be expected to lie above the simulated background. In this sense, GDAS profiles can still be used to detect clouds in data collected with our cloud cameras.

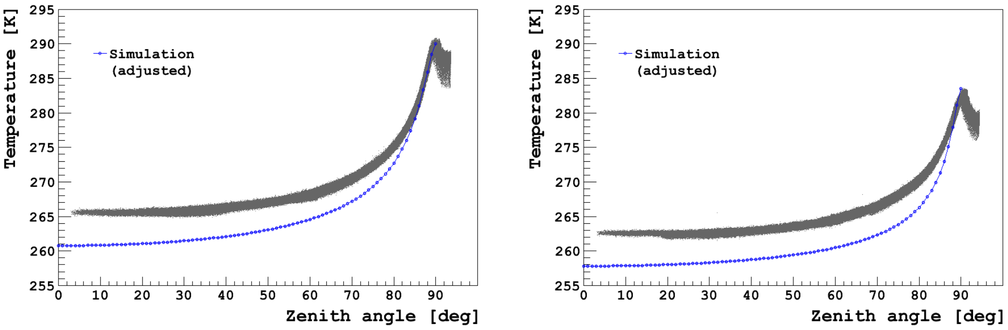

(a)

(b)

Figure 5.22: Measurements made by the Los Leones infrared cloud camera on two different clear nights. **(a)** is the same camera data as in Figure 5.21a and **(b)** is the same camera data as in Figure 5.21b. The effective sky temperatures have been simulated using MODTRAN with corresponding GDAS atmospheric profiles. The GDAS temperature lapse rates from the surface (at 1420 m a.s.l.) up to 5 km a.s.l. have been adjusted such that ground temperatures agree with cloud camera observations (sky temperature measurements at the horizon).

Chapter 6

Estimations of Cloud Base Heights

When analysing extensive air showers at the Pierre Auger Observatory, knowledge of cloud positions above the array are important (see Section 1.2.3). Although our infrared cameras can provide the angular positions of clouds within the fluorescence telescopes' fields of view, they cannot directly measure their heights. On the other hand, the lidars, the Central Laser Facility (CLF), and the eXtreme Laser Facility (XLF) can determine cloud heights, but not their angular coordinates. The various instruments complement one another and together provide useful geometry for air shower reconstruction.

In this chapter, I explore two methods for estimating cloud base heights at Auger. The first study uses sky temperature measurements made by a single-pixel infrared radiometer. The second study predicts cloud layers using atmospheric temperature and humidity profiles obtained from the Global Data Assimilation System (GDAS). This research could potentially offer supplementary techniques for retrieving cloud height information at the Observatory. In the event that other instruments are non-operational, these methods may be appropriate substitutes.

6.1 Determining cloud base heights using a single-pixel radiometer

We have discussed that the majority of water clouds (such as stratus) appear bright in the night sky at infrared wavelengths (see Section 2.4). As the emission spectra of water clouds resemble those of a black body, the intensity of their thermal radiation strongly corresponds to their temperatures. This, in turn, depends on the clouds' heights in the atmosphere. It follows that measuring the temperature of a cloud can be used to provide some indication of its altitude. As such clouds are opaque at infrared wavelengths, a ground-based observer measures the cloud base height.

A number of past studies have determined cloud base heights using this principle [45][104][105][106]. Common examples of instruments that detect thermal emission from clouds are infrared cameras and radiometers, which have been discussed in Chapter 3. Measurements of infrared brightness from clouds can be converted into temperatures, and heights can then be estimated. This can typically be verified with a lidar or ceilometer.

In this section, I attempt to retrieve cloud height information at Auger from the radiometer installed at the Central Laser Facility (CLF). As discussed in Section 3.5, the CLF operates during fluorescence detector (FD) data acquisition to provide hourly measurements of the minimum cloud base heights. This provides the truth for my study. Although a similar method could be applied by comparing cloud temperatures measured with our infrared cameras to heights observed by the lidars, the downtime of the lidars, or the unavailability of lidar data in more recent years, meant this was not possible. Moreover, at earlier times when there were overlapping measurements, the temperature calibrations developed for our cameras could not be used due to the different camera settings. For these reasons, the most straightforward approach was to use data from the single-pixel radiometer co-located with the CLF. A previous study that used the same model of radiometer to determine cloud base heights has formed the basis of this research [45].

6.1.1 Data selection

Our single-pixel radiometer is mounted so that it records an overhead sky temperature within the 5.5–14 µm spectral range of the detector. As discussed in Section 3.2, the sky temperatures are determined by measuring the temperature difference between its 3° field of view and a reference. This difference is known as the uncompensated temperature, and the reference is the local temperature of the detector canister. As the canister is fixed near the ground, the reference approximates surface temperatures. Within the band-pass of our detector, clear sky uncompensated temperatures are typically about $-20\,^{\circ}\mathrm{C}$, but appear warmer when there is more humidity in the air. Given the uncompensated temperature can provide the temperature difference between a cloud in the radiometer's field of view and the ground, it is appropriate to compare this measurement to the minimum cloud base heights observed with the CLF.

The radiometer performs measurements at 5 minute intervals that can be compared to the hourly minimum cloud base heights reported by the CLF. As discussed in Section 3.5, the positions of clouds that lie between the laser beam and an FD building cannot be known from this bistatic lidar technique. The ambiguous geometry means that any reported cloud heights do not reflect their true values. In contrast, when clouds are directly above the laser, their heights can be determined. The existing analysis does not distinguish between both cases. However, an improved cloud detection algorithm that can is currently being developed by Violet Harvey, a PhD. student at the University of Adelaide. I have proceeded with my study using the minimum cloud base heights determined from the existing CLF analysis. To remove the ambiguity in the clouds' geometries in my study, preliminary results from the new analysis have instead been used to identify when clouds were directly above the CLF. Given the small field of view of our radiometer, it is likely that both instruments view the same cloud. To ensure the radiometer measurements are coincident with the hourly CLF data, only observations corresponding to the maximum uncompensated temperature within each hour are selected. This provides the best match to the minimum cloud base heights recorded. The data covers all of 2010 and the results are presented in Figure 6.1.

6.1.2 Results

The measured sky temperatures are influenced by thermal emission from clouds, and to a lesser extent, atmospheric water vapour. Figure 6.1 suggests that the sky temperatures gradually cool as clouds are found at higher altitudes. As clouds are at similar absolute temperatures to the air that surrounds them, such a relationship can be expected. However, the temperature lapse rate of $1.6\,^{\circ}\mathrm{C\,km^{-1}}$, determined by a fit to the data, appears to be quite low. Using predictions made by the GDAS model, the mean environmental lapse rate for these data is found to be $6.9\,^{\circ}\mathrm{C\,km^{-1}}$. Consequently, the sky temperatures are much warmer than expected.

One possible source for the apparent discrepancy may be due to water vapour between the clouds and the detector. Although atmospheric attenuation at these wavelengths is low, intervening water vapour can affect ground-based measurements, as demonstrated in Section 2.3. Not only does water vapour emission add to the infrared flux observed at the ground, but it also obscures downwelling thermal radiation emitted by clouds [99]. The extent of both effects depends on the amount of water vapour within the radiometer's line-of-sight. For these data, the amounts of precipitable water vapour were calculated with GDAS and found to range from 0–15 mm. When comparing between sky temperatures on dry (taken as $\mathrm{TPWV} < 6\,\mathrm{mm}$) and more

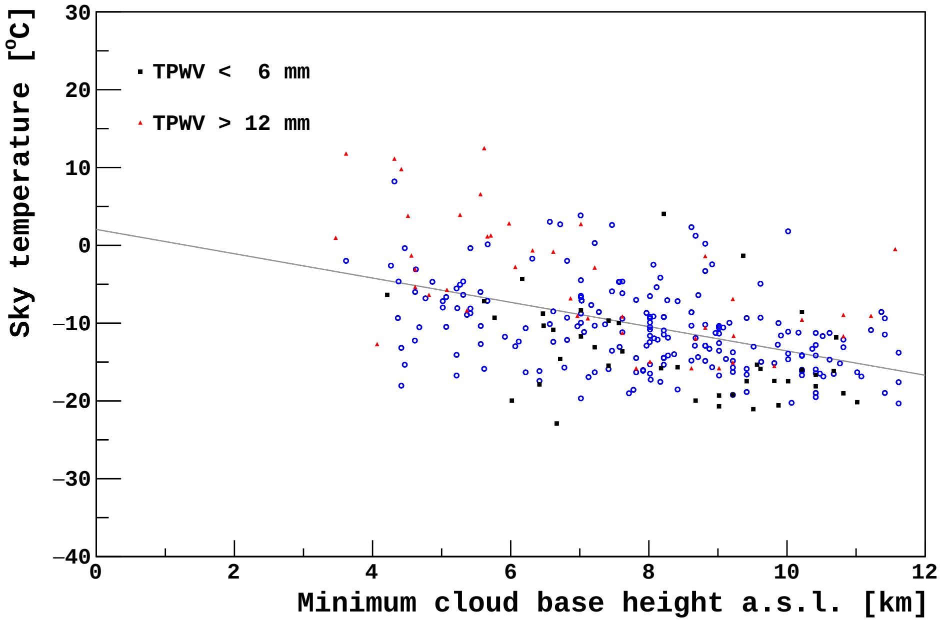

Figure 6.1: A comparison between sky temperatures measured using the radiometer and cloud base heights observed with the Central Laser Facility (CLF) at the Pierre Auger Observatory. The CLF provides hourly measurements of the minimum cloud base heights. A cut has been applied to identify clouds that were directly above the CLF. The fitted line uses all of the data, and indicates a lapse rate of 1.6 °C km^{-1}. The residual effects on sky temperatures due to water vapour can also be seen, with warmer temperatures generally corresponding to higher humidities. Humidities are expressed in units of total precipitable water vapour (TPWV), the depth of water in an atmospheric column, if all the water precipitated as rain. Data covers all of 2010, and ground-level is at 1401 m a.s.l.

humid (TPWV > 12 mm) nights, the effects of water vapour can be seen. Figure 6.1 shows that additional water vapour generally results in warmer temperatures. However, regardless of the amount of moisture, water vapour only appears to have a small effect on the measured sky temperatures. Given that atmospheric absorption and emission are reduced within the band-pass of the radiometer, the infrared flux at ground is expected to be largely thermal radiation emitted by clouds. To this end, the radiative properties of the clouds in this study are investigated.

The majority of observations are likely to be of ice clouds (such as cirrus), given the heights at which they around found. While thermal emission from water clouds approximates black body radiation, the emission spectra of ice clouds have been shown to be rather weak (see Section 2.4). This property is described by their emissivities, which are characteristically much lower than for water clouds. To the extent that cirrus clouds do not behave as a black body, they are then reflective (conservation of energy). Hence, the infrared flux detected by the radiometer would mostly be a combination of thermal emission from a cloud and some reflected component of thermal radiation emitted by the ground. A simple model used in [45] to describe this behaviour is:

$$\sigma T_{sky}^4 = \epsilon_{cloud}\sigma T_{cloud}^4 + \epsilon_{ground}\sigma T_{ground}^4(1 - \epsilon_{cloud}) \,, \tag{6.1}$$

where the sky temperature T_{sky}, cloud base temperature T_{cloud}, and ground temperature T_{ground} are in Kelvin, and σ is the Stefan-Boltzmann constant. Thermal radiation is emitted by the cloud with an emissivity ϵ_{cloud}, and the ground with an emissivity ϵ_{ground}.

Clouds are at similar absolute temperatures to the air that surrounds them. As a result, cloud base temperatures can be directly determined from atmospheric temperature profiles, such as those measured with a radiosonde or, in my case, predicted by the GDAS model. An alternative approach is to infer atmospheric temperatures from the environmental lapse rate and the temperature at the ground. In Section 3.6.3, comparisons between the GDAS temperature data and measurements made by sondes, and weather stations, showed some discrepancies. Given that lapse rates are a differential temperature measurement, they would be less affected by systematic errors in the GDAS model. For this reason, I chose to infer cloud base temperatures using the latter method. The temperature of a cloud base T_{cloud} at a height h in the atmosphere can then be calculated from:

$$T_{cloud} = T_{ground} - \Gamma \cdot h \,, \tag{6.2}$$

where the environmental lapse rate Γ is determined from a fit to the GDAS temperature profile, and the ground temperature T_{ground} is measured by the radiometer's internal thermistor.

It follows that Equation 6.1 and Equation 6.2 can be used to estimate a cloud base height. The additional quantities that must be known are the emissivities of both the ground and the cloud. I estimate the ground emissivity at the Pierre Auger Observatory to be 0.96 ± 0.02 within our spectral band [1]. Since the cloud base heights in this study are already known,

---

[1] Based on the 10.8 µm channel from the CIMSS IREMIS Global Infrared Land Surface Emissivity: UW-Madison Baseline Fit Emissivity Database [107]. Data is available online at `http://cimss.ssec.wisc.edu/iremis/`.

the equations can instead be rearranged to determine a typical range of emissivities for the clouds seen above the array. The results are presented in Figure 6.2. The cloud emissivities are found to have a mean of 0.30 and a standard deviation of 0.12. These values seem reasonable and are in agreement with those expected for cirrus clouds [28][43][108][109]. At lower altitudes the emissivities would be expected to be closer to 1. However, the existing cloud detection algorithm used by the CLF can only detect clouds that are at least 2 km above the ground, so this could not be verified.

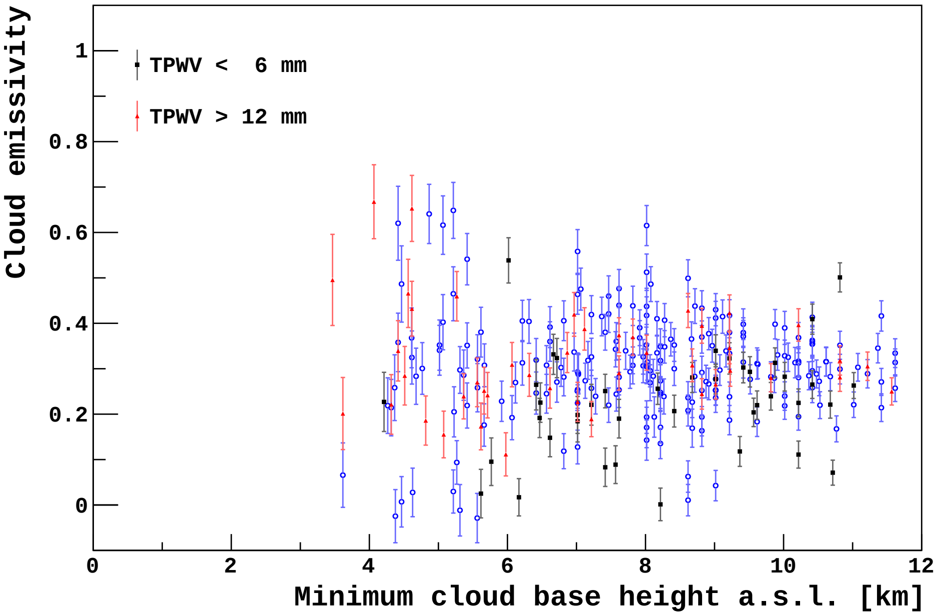

Figure 6.2: The predicted emissivities for the clouds in this study. The emissivities have a mean value of 0.30 and a standard deviation of 0.12. Variations in humidity can be seen to have a negligible effect on these calculations. Humidities are expressed in units of total precipitable water vapour (TPWV), the depth of water in an atmospheric column, if all the water precipitated as rain. Error bars derive from the uncertainty in the ground emissivity.

Using the average cloud emissivity of 0.30 that I have calculated, some estimates can now be made on the cloud base heights. The predicted heights minus those recorded by the CLF are shown in Figure 6.3. The distribution has a mean of 0.5 km and a standard deviation of 4.1 km. The large uncertainties in the heights determined with this method suggest that it would be a poor substitute for the instruments currently operating at Auger.

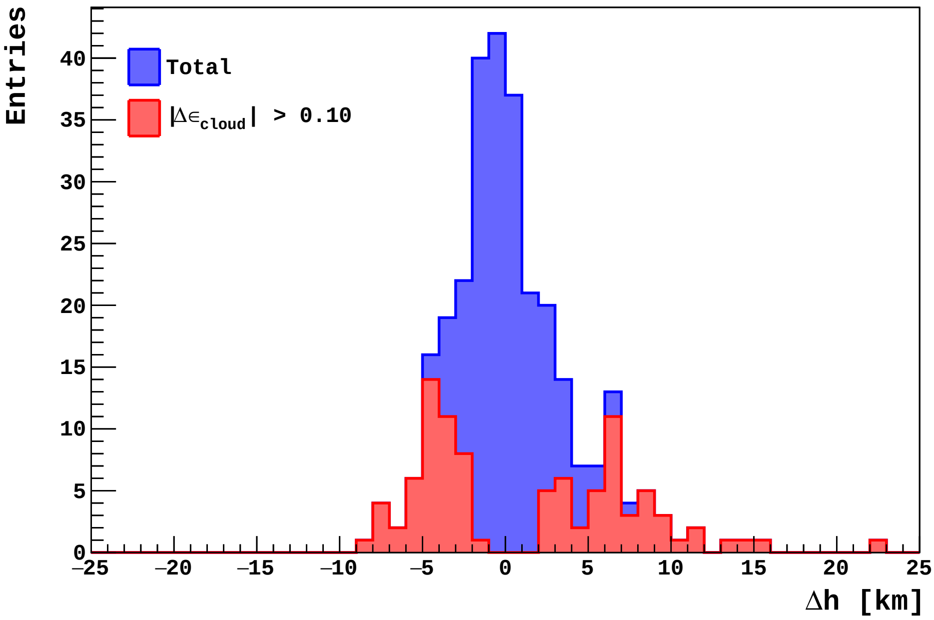

Figure 6.3: Distributions of the cloud base heights derived with the radiometer minus the true heights recorded with the CLF. There is poor agreement when the estimated cloud emissivities deviate from the mean value of 0.30. Those that differ by more than 0.10 are in red.

In Figure 6.3, the disagreements are greatest for clouds with estimated emissivities (from Figure 6.2) that have deviated significantly from the mean value (those that differ from the mean by more than 0.10 are indicated). The sensitive nature of the estimated cloud base heights on cloud emissivity is not surprising considering that the value determines both the intensity of thermal radiation emitted by the cloud, in addition to the amount of upwelling radiation reflected at the cloud base. This response can be understood by calculating the partial derivative of the cloud base height h with respect to the cloud emissivity ϵ_{cloud}, which is given by:

$$\frac{\partial h}{\partial \epsilon_{cloud}} = \frac{T_{sky}^4 - \epsilon_{ground} T_{ground}^4}{4 \Gamma \epsilon_{cloud}^2} \cdot \left(\frac{T_{sky}^4 - \epsilon_{ground} T_{ground}^4 (1 - \epsilon_{cloud})}{\epsilon_{cloud}} \right)^{-\frac{3}{4}} . \quad (6.3)$$

Substituting in nominal values ($T_{sky} = 260\,\mathrm{K}$, $T_{ground} = 280\,\mathrm{K}$, $\epsilon_{ground} = 0.96$, and $\Gamma = 6.9\,°\mathrm{C\,km^{-1}}$), it can be shown that:

$$\left. \frac{\partial h}{\partial \epsilon_{cloud}} \right|_{\epsilon_{cloud}=0.30} \simeq -70\,\mathrm{km} . \quad (6.4)$$

Consequently, the cloud emissivities that deviate from the mean by 0.10 can be expected to result in cloud base heights that are incorrect by $\sim$ 7 km. This is similar to the errors observed in Figure 6.3.

6.1.3 Conclusions

This study set out to determine whether the single-pixel radiometers at the Pierre Auger Observatory would be able to retrieve cloud base height information. One of the radiometers is installed at the CLF and records the overhead sky temperatures at 5 minute intervals. These temperatures were found to correlate with the cloud base heights measured by the CLF, with cooler sky temperatures found when the clouds were at higher altitudes. However, the sky temperatures are much warmer than would be expected, assuming a typical environmental lapse rate.

A likely explanation is that clouds at those heights not only emit thermal radiation, but also reflect some fraction of the Earth's upwelling radiation. Both depend on the emissivities of the clouds. With the simple model used in [45], I have calculated the emissivities of the clouds in this study to have a mean of 0.30 and a standard deviation of 0.12. These results are consistent with [45] and also agree with previous experiments that determined the emissivities of cirrus clouds.

Assuming a cloud emissivity of 0.30 and using radiometer measurements, the cloud base heights are estimated and compared to the true heights recorded by the CLF. Although the average height difference of 0.5 km is encouraging, the large uncertainties as a result of this method are rather disappointing. This can be explained by the sensitivity that predicted heights have on the cloud emissivity. The resultant height differences have a standard deviation of 4.1 km. This is greater than the spread of $\sim$ 700 m found in [45], however, the clouds in this study are at much higher altitudes, so their thermal emissions are more difficult to distinguish from a clear sky (refer to Figure 2.11).

Given the above, it can be concluded that the cloud base heights determined by the single-pixel radiometers using this method, are not precise enough to replace the instruments currently operating at Auger. This research, however, would benefit from CLF measurements of low level clouds. As clouds at lower altitudes approximate black bodies, a more direct relationship between the radiometer sky temperatures and cloud base heights can be expected.

6.2 Predicting cloud layers using GDAS atmospheric profiles

Other than inferring the heights of clouds from the thermal radiation they emit, there are various techniques that instead use atmospheric parameters to predict cloud layers. They range from simple practices that use surface level data [110] to more direct methods that use upper air measurements from radiosondes [111][112][113]. A common feature, however, is to relate the presence of clouds to the temperature and moisture content of the air.

In this section, I attempt to retrieve cloud height information at the Pierre Auger Observatory from temperature and humidity profiles predicted by the Global Data Assimilation System (GDAS). This study is based on a method that was developed to predict the boundaries of cloud layers from radiosonde launches [113].

6.2.1 Method

Given that clouds are regions of moist air, a higher relative humidity can be expected within a cloud layer than outside of its boundaries. Not only is the humidity expected to change, but also the temperature of the atmosphere. The flux exchange between the cloud top and the layer of air above acts to produce cooling at the top of the cloud [30]. In contrast, the infrared radiation trapped between the Earth's surface and the cloud base [30], and the release of latent heat near the base of the cloud [32], warms the surrounding air. Additionally, temperature inversions provide favourable conditions for layered clouds to form [113]. For these reasons, there is an expectation that either relative humidity or temperature could be used to predict the presence of clouds.

In this study, clouds are characterised in the same manner as [113], by inflections in both their relative humidity $R(h)$ and temperature $T(h)$ profiles. The criteria that each cloud layer must satisfy are:

$$R''(h) \leq 0 \quad \text{and} \quad T''(h) \geq 0\,, \tag{6.5}$$

where $R''(h)$ and $T''(h)$ are the second derivatives of the relative humidity and temperature vertical profiles, with respect to height. The magnitudes of these quantities are not important [113].

In addition to predicting cloud boundaries, that analysis was supplemented with a separate method to determine fractional cloud cover (or cloud amount) [114]. The technique relates radiosonde air temperature and dew point depression (air temperature minus the dew point) to cloud amount,

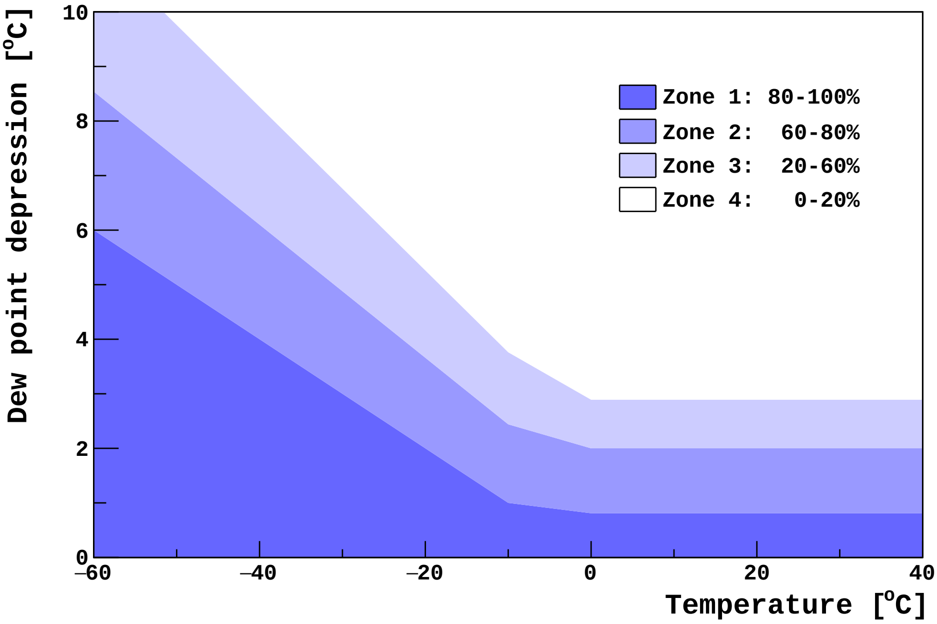

Figure 6.4: The Arabey diagram, which is used to predict cloud cover from radiosonde temperature and dew point depression (air temperature minus the dew point) measurements [114]. The cloud cover (or cloud amount) refers to the fraction of the sky obscured by clouds. Each zone represents a different amount of cloud cover. Zone 1 is an area of complete saturation with a cloud cover of 80–100 %, Zone 2 is an area of near saturation with a cloud cover of 60–80 %, Zone 3 is an area of partial saturation with a cloud cover of 20–60 %, and Zone 4 is an area of dry air with a cloud cover of 0–20 %. Image modified by the author from [113].

and can be summarised graphically with the Arabey diagram (see Figure 6.4). The diagram contains four zones that correspond to different levels of saturation. The zones range from completely saturated air with 80–100 % cloud cover (zone 1) to dry air with 0–20 % cloud cover (zone 4). In this analysis, the minimum dew point depression within a cloud layer, and the corresponding temperature, are used to estimate the cloud cover from the Arabey diagram.

Although the main focus of my research has been to determine cloud base heights from the GDAS data, information on fractional cloud cover could potentially be estimated using this technique. Each of the Auger lidar stations report an hourly minimum cloud base height along with the fractional cloud cover, as determined by their scanning routines (see Section 3.4 for details).

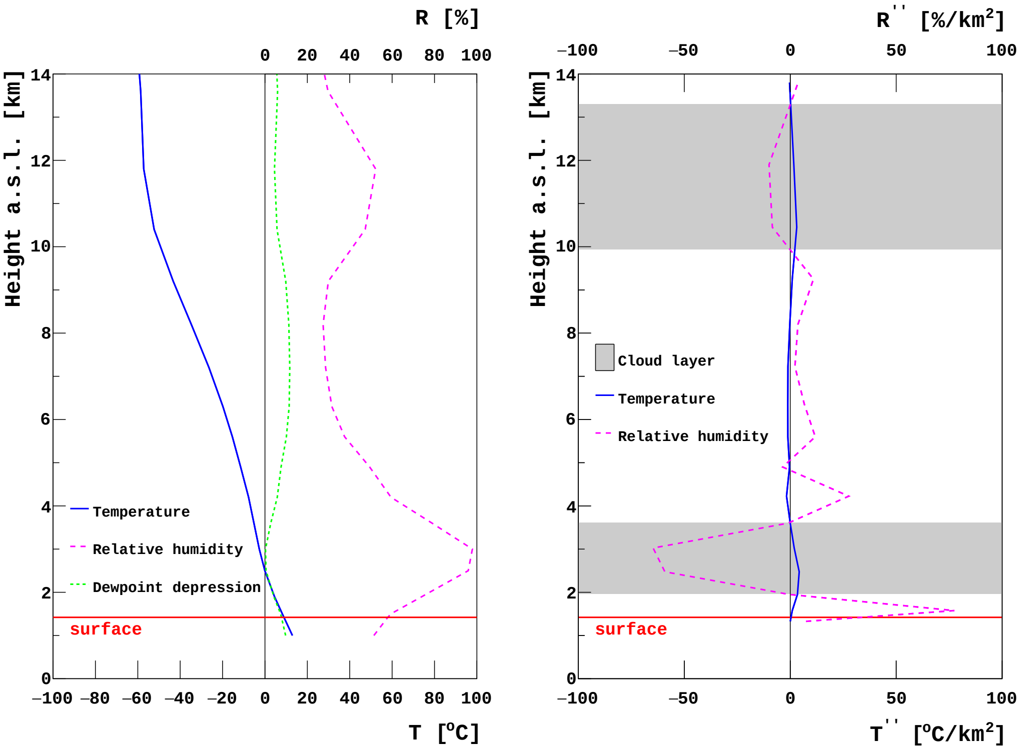

(a) (b)

Figure 6.5: GDAS atmospheric profiles for a cloudy night on 2012-10-10 06:00:00 UTC. **(a)** Temperature, relative humidity, and dew point depression profiles. **(b)** The second derivatives of the temperature and relative humidity profiles. Possible cloud layers satisfy the criteria in Equation 6.5 and are shaded grey. The Los Leones lidar station observed an hourly minimum cloud base height at 2551 m a.s.l. and the corresponding cloud top to be 3211 m a.s.l. The horizontal line at 1420 m a.s.l. indicates the surface level at Los Leones.

To test the applicability of this method to the Observatory, the lidars, which are installed at each of the main fluorescence detector (FD) sites, provided the truths for my study. Given that only a single GDAS grid point (refer to Figure 3.14) is used to describe the atmospheric conditions above the entire array, a separate analysis is performed for each lidar. To make a more direct comparison between the lidar observations and the GDAS model, only the lidar data coincident with the 3-hourly model are used. The upper atmospheric quantities predicted by the model are available at 23 constant pressure levels and have been listed in Table 3.1. However, the profiles that

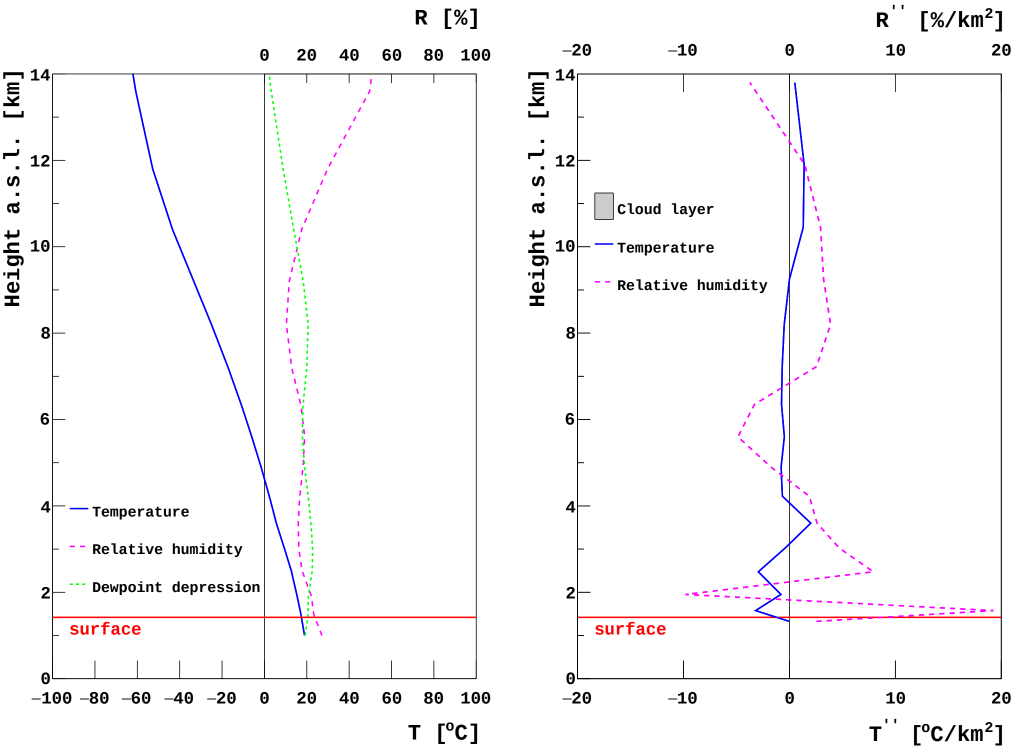

(a) (b)

Figure 6.6: GDAS atmospheric profiles for a clear night on 2014-01-04 03:00:00 UTC. **(a)** Temperature, relative humidity, and dew point depression profiles. **(b)** The second derivatives of the temperature and relative humidity profiles. No cloud layers are predicted for this atmosphere, and none were observed by the lidars. The horizontal line at 1420 m a.s.l. indicates the surface level at Los Leones.

are stored within the Auger $\overline{\text{Off}}\underline{\text{line}}$ database have been interpolated between those fixed levels. When using Equation 6.5 to predict cloud layers contained within the GDAS profiles, it was found that calculating the second derivatives $R''(h)$ from $R(h)$, and $T''(h)$ from $T(h)$, using step sizes Δh that were too small, resulted in falsely identified cloud layers. For instance, a single cloud layer would instead be divided into multiple thin layers. The best results were achieved when the approximate height resolution of the GDAS model was preserved. The heights of the 1976 U.S. Standard Atmosphere that correspond to the 23 fixed GDAS pressure levels have been used in my analysis, and can also be found listed in Table 3.1. As water vapour pressure is the quantity that is stored in the database, conversions into both relative humidity and dew point temperature were made for this study. The details of the conversions are discussed in Section 2.3.

To illustrate how the method identifies potential cloud layers within the GDAS profiles, two examples are shown here; the first is a sky that was identified to be overcast by the lidars, and the second was observed to be clear. Figure 6.5a shows atmospheric temperature, relative humidity, and dew point depression with height. An increase in the relative humidity at $\sim$ 3 km a.s.l. and $\sim$ 12 km a.s.l. corresponds to a decrease in the dew point depression (i.e. the air temperature approaches the dew point). This implies there is more moisture in the atmosphere, and an initial expectation is that clouds may be found at those heights. Figure 6.5b shows the second derivatives of the temperature and relative humidity profiles. The two shaded regions satisfy the criteria in Equation 6.5 and indicate possible cloud layers. The lowest layer has the cloud base at 1960 m a.s.l. and the cloud top at 3612 m a.s.l. The highest layer has the cloud base at 9940 m a.s.l. and the cloud top at 13 300 m a.s.l. Both layers have a cloud cover of 100 %, which were calculated from the Arabey diagram (Figure 6.4). The Los Leones lidar station observed an hourly minimum cloud base at 2551 m a.s.l. and the corresponding cloud top to be at 3211 m a.s.l. The cloud cover measured by the lidar was found to be 100 %. In this example there is good agreement between the minimum cloud base height predicted from the GDAS data and the correct height reported by the lidar. In fact, the true cloud layer is contained within the cloud boundaries predicted using this method. Not only do the cloud heights agree, but also the cloud cover. The overall agreements with each of the lidars will be discussed in the results.

Figure 6.6 shows the same procedure applied to a sky that was observed by the lidars to be completely clear. In Figure 6.6a the relative humidities at most heights are low (or the dew point depressions are high), indicating the atmosphere is drier than in the previous example. As a consequence, there is an expectation that clouds are less likely to be found for those atmospheric conditions. The second derivatives of the temperature and relative humidity profiles are shown in Figure 6.6b. However, unlike the previous demonstration, there are no regions that satisfy the criteria. As a result, no clouds layers are predicted for that atmosphere.

A simple requirement that no cloud layers are found in order to identify a clear sky is too harsh in practice. In this study, the same indicator has been adopted as [113] to determine whether a sky is clear. If a cloud layer has a cloud cover $\leq 20\,\%$ (i.e. zone 4 on the Arabey diagram) then it is considered to be clear. Similarly, when the cloud cover is greater than 20 %, but with a layer thickness ≤ 300 m, it is also said to be clear. Only when all cloud layers are clear, according to these rules, is the atmosphere diagnosed as a clear sky. In contrast, a cloudy sky is when at least one layer has a cloud cover greater than 20 % and a thickness greater than 300 m.

6.2.2 Results

In order to assess the applicability of the method to the GDAS data, comparisons have been made to observations from the four lidar stations operating at Auger. The first set of analyses compares times that were identified as clear by each lidar to predictions based on the atmospheric profiles. In passing, clouds were often found in the GDAS profiles above heights of 8 km a.s.l., even when the lidars recorded clear skies. A likely explanation is due to the high relative humidities frequently found within the profiles at those altitudes, rather than a deficiency of the method. I suspect that relative humidities at those heights are incorrectly calculated (see the Magnus formula in Equation 2.9) when the air temperatures are very low $\sim -40\,^{\circ}\mathrm{C}$. For the purpose of this study, lidar observations are considered to be clear when either the cloud cover is recorded as 0 % or the minimum cloud base height is greater than 8 km a.s.l. Likewise, the GDAS profiles are considered to be clear when no cloud layers, according to the rules that were outlined in the method, are found below 8 km a.s.l.

Table 6.1 presents a summary of successful clear sky predictions P_{clear} for each lidar station, determined from N observations. The findings show the overall probability of successfully predicting a clear sky is $\sim 80\,\%$. In addition, the agreement appears to be largely independent of location. Taken together, these results suggest the information contained within the GDAS data could be used to test whether the sky is likely to be clear above the array.

Site	$\boldsymbol{P_{clear}}$ [%]	$\boldsymbol{N}$
Los Leones	79.9	1178
Los Morados	76.0	508
Loma Amarilla	77.0	1489
Coihueco	79.3	1106

Table 6.1: Frequencies of correct clear sky predictions P_{clear} at each Auger fluorescence detector site, based on N lidar observations. Data covers the period of 2009-01-01 to 2014-03-10.

Given the promising results at predicting clear skies, the next test is to investigate the performance when clouds are present. In this study, cloudy skies are when a lidar observes a cloud cover greater than 0 % and the minimum cloud base height is below 8 km a.s.l. Table 6.2 provides a breakdown of successful cloudy sky predictions P_{cloud} for each lidar station. The success at estimating the minimum cloud base height P_{height}, in addition to the cloud cover $P_{height+cover}$ are also shown. Since the atmospheric profiles measure the level at which condensation occurs, whilst the lidars detect clouds when the concentration of cloud particles surpasses some limiting value, the cloud heights are not expected to be in perfect agreement [113]. When the true minimum cloud base height is contained within the cloud boundaries predicted for the lowest layer (for example the bottommost shaded region in Figure 6.5b), I consider that height to be successfully calculated. The cloud cover of that layer is correct when the lidar observation lies within the range predicted from the Arabey diagram.

Site	$\boldsymbol{P_{cloud}}$ [%]	$\boldsymbol{P_{height}}$ [%]	$\boldsymbol{P_{height+cover}}$ [%]	$\boldsymbol{N}$
Los Leones	60.1	27.0	10.5	715
Los Morados	53.5	27.0	13.0	185
Loma Amarilla	66.1	24.1	8.7	623
Coihueco	51.9	19.0	6.2	759

Table 6.2: Frequencies of correctly predicting a cloudy sky P_{cloud} at each Auger fluorescence detector site, based on N lidar observations. The success at predicting the minimum cloud base height P_{height}, in addition to the cloud cover $P_{height+cover}$ are also shown.

It is apparent from both tables that the success at predicting a cloudy sky is lower than for a clear sky. Moreover, the cloud heights are correct only ~ 25 % of the time. Given that the main goal is to determine cloud base heights from the GDAS data, these results are rather disappointing. However, there are several possible reasons for these disagreements that will be discussed in the conclusions.

Although the overall success at predicting the cloud base heights is quite low, there are also some encouraging results. Table 6.3 shows a selection of events that have been chosen at times when the cloud cover was high (both the lidar cloud cover and predicted cloud cover are greater than 80 %). The findings suggest that cloud heights are more accurately predicted when the

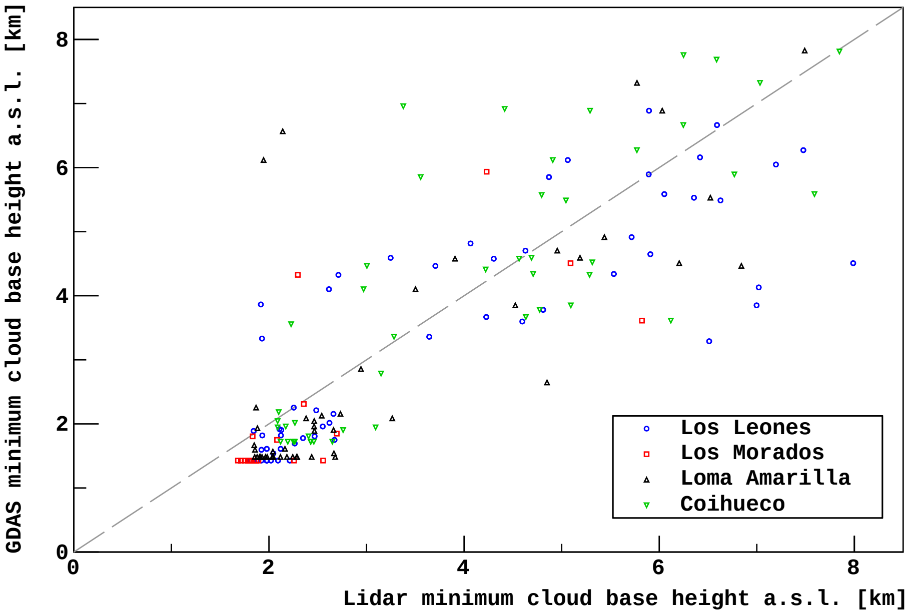

Figure 6.7: The minimum cloud base heights predicted from the GDAS profiles compared to the minimum values recorded by the lidars. A one-to-one line is indicated by the dashed line. Only times when the lidar cloud cover and the predicted cloud cover are greater than 80 % are shown.

sky is overcast, and worse when there is broken cloud. This is not surprising considering a single GDAS profile describes the atmospheric conditions above the entire array. The best agreements are expected to be at times that are either totally clear or completely overcast. At the times when the cloud cover was high, the accuracy of the cloud base heights has been investigated. Figure 6.7 shows comparisons between the true minimum cloud base heights and those predicted from the GDAS data. It can be seen there is reasonable agreement, with both values roughly following a one-to-one line. The standard deviation of all the height differences is ~ 1.1 km. The spread is reduced slightly at lower altitudes. Possible reasons for the spread and disagreements are discussed presently.

6.2.3 Discussion and Conclusions

The initial goal of this study was to investigate whether atmospheric profiles from the GDAS model could be used to provide cloud height information for the Pierre Auger Observatory. For this analysis the cloud layers were

Site	P_{height} [%]	N
Los Leones	66.7	63
Los Morados	84.2	19
Loma Amarilla	68.1	47
Coihueco	50.0	46

Table 6.3: Frequencies of correctly predicting the minimum cloud base height P_{height} at each Auger fluorescence detector site from N lidar observations. A subset of Table 6.2, only showing the times when both the lidar cloud cover and the predicted cloud cover were greater than 80 %.

predicted using the same criteria as [113], a technique that was originally used to determine cloud heights from radiosonde temperature and relative humidity data. In addition, that analysis was supplemented by a method to determine cloud cover using measurements of temperature and dew point depression.

To test the performance of the method when applied to the GDAS data, the four lidar stations operating at the Observatory provided the truths for my study. Even at times observed to be clear by the lidars, it was often found that clouds were predicted at heights above 8 km a.s.l. This is most likely due to the high relative humidities found within the GDAS profiles at those heights rather than the method. Below 8 km a.s.l. the performance has been investigated.

Clear skies were found to be successfully predicted about 80 % of the time. The success at predicting cloudy skies, however, is significantly worse at 60 %. In addition, only 25 % of the minimum cloud base heights agree with the lidars, whilst just 10 % have the correct heights *and* cloud cover. One contributing factor to these poor agreements is likely the 1° × 1° latitude-longitude spatial resolution of the GDAS profiles, which are used to describe atmospheric conditions for the entire array. It seems reasonable that the best agreements would occur at times when the sky is either clear or overcast. In fact, when multiple lidar stations report a clear (or overcast) sky, then the sky cover is approximately uniform across the array, and there is better agreement with the predictions. This can also be seen when there are high levels of cloud cover above a single lidar, as opposed to broken layers of cloud. When the cloud cover is high, around 70 % of the minimum cloud base heights are correctly predicted at the Los Leones and Loma Amarilla

FD sites. These agreements are slightly worse at Coihueco, and seemingly the best at Los Morados, although with fewer lidar observations. An initial expectation was that the best results would be found at Loma Amarilla considering it is located nearest to the GDAS grid point (refer to Figure 3.14). However, no evidence in this study has been found to support this.

Apart from the spatial resolution of the GDAS data, the 3-hourly temporal resolution of the model is also expected to introduce some uncertainties into these predictions. Although the molecular atmosphere is very stable over 3-hour periods [21], the presence of clouds above the array is known to vary over much shorter time scales. These dynamics can be seen in the images captured with the infrared and optical cloud cameras, and by the GOES weather satellite, in addition to data from the single-pixel radiometers, lidars, and the CLF/XLF. Moreover, these dynamics may not be accurately described by the global model, particularly the influence the neighbouring Andes mountains may have on the climate above the array [21] (possible examples are in Section 6.3). A more abstract source for these discrepancies is that although atmospheric conditions may be favourable for cloud, there may be reasons that cloud formation does not occur. For instance, a lack of cloud condensation nuclei onto which water vapour can condense. Nevertheless, whether or not cloud layers are predicted is ultimately constrained by the height resolution of the model (the 23 constant pressure levels listed in Table 3.1). This is the most likely reason for the low success at finding clouds that are observed. At many heights the resolution of the atmospheric temperature and humidity profiles is only $\sim 0.5\,\mathrm{km}$. With a limited resolution, some information is lost between those fixed pressure levels.

Although this study set out to determine the heights of clouds from the GDAS data, the current results show it to be unreliable for most purposes. However, there does appear to be useful information contained within the GDAS profiles. In particular, the profiles could serve to identify whether the sky is likely to be clear or cloudy. Separate studies, with these same goals, have been investigated by Jarryd Day, a current student at the University of Adelaide [115]. In that analysis, estimates on clear or cloudy conditions had a similar accuracy of $\sim 70\,\%$, with preliminary results showing promising results of retrieving cloud base heights.

Given that the GDAS data is already made available to us within the $\overline{\text{Off}}\underline{\text{line}}$ database [21], it may be useful to incorporate a routine to test the likelihood of a sky being clear or cloudy, based on the atmospheric temperature and humidity profiles at that time. This information could potentially supplement data from the other instruments operating at the Observatory. An example of a routine may be found in Appendix D.

6.3 Concluding remarks

One topic of interest is the influence that the neighbouring Andes mountains can have on the climate at the Pierre Auger Observatory. In relation to my work, I was interested in what impact the terrain could have on the population of clouds above the array. For instance, cloud dynamics are known to be influenced by geographical features (such as mountains [36][116]). Given the extensive network of instruments used to detect clouds at the Observatory (see Chapter 3), there are several ways to investigate the cloud populations at Auger. In this section, I briefly summarise some results acquired using data from infrared satellite images of the Observatory (see GOES, Section 3.3), and measurements made by the lidars located at the main fluorescence detector sites (see Section 3.4).

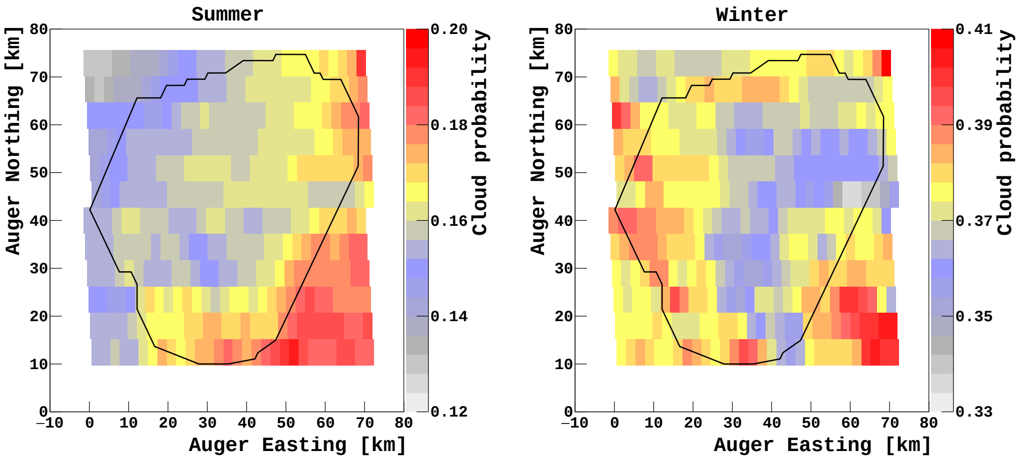

(a) (b)

Figure 6.8: Mean cloud probabilities at the Pierre Auger Observatory using infrared satellite data (see GOES, Section 3.3). Seasonal averages are shown for **(a)** Summer months and **(b)** Winter months. Note that the colour scales are different in order to highlight the excess of clouds seen to the south-east of the Observatory. An outline of the Observatory is in black. Data covers 2009 through to 2015. Illustrated by the author of this dissertation.

Figure 6.8 shows cloud probabilities across Auger that have been averaged during Summer (Dec–Feb) and Winter (Jun–Aug) months. One feature that stands out is the apparent excess of clouds towards the south-east of the array. A gradient in the cloud cover had also been found in [60]. Although the overall cloud cover is greatest during Winter, the excess appears stronger in Summer. Interestingly, the excess is located in the vicinity of

Llancanelo Lake. The environment could perhaps provide favourable conditions for cloud formation, possible due to greater amounts of water vapour in the area. Another explanation may be that the wetlands have different surface emissivities. Depending on the emissivity in the satellite's infrared channels, this could lead to false cloud detections (using the algorithm discussed in Section 3.3.2).

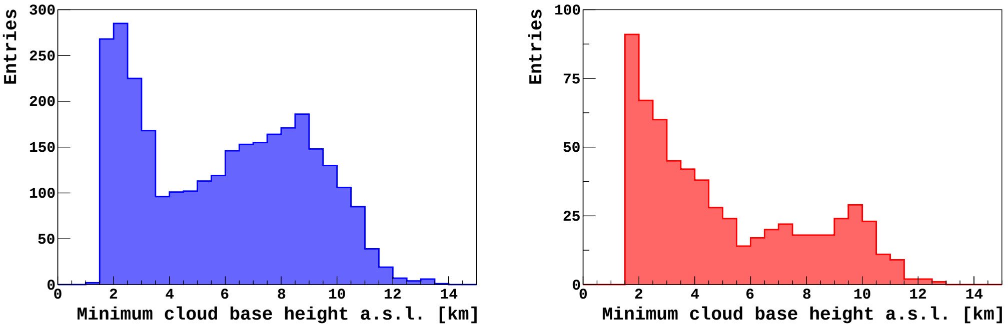

(a) Los Leones **(b)** Los Morados

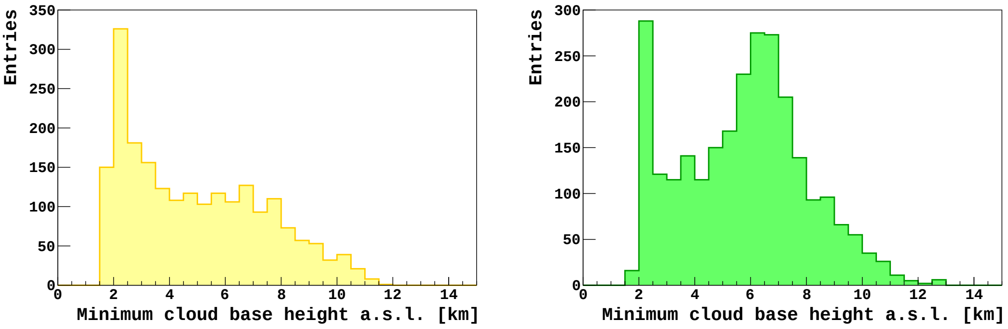

(c) Loma Amarilla **(d)** Coihueco

Figure 6.9: Distributions of minimum cloud base heights recorded with the lidars at the Pierre Auger Observatory, which are installed at the four main fluorescence detector sites. An excess can be seen with the Los Leones lidar at about 8–10 km a.s.l., and with the Coihueco lidar at about 6–8 km a.s.l. Data covers June 1st, 2009 through to March 10th, 2014. Illustrated by the author of this dissertation.

Figure 6.9 presents distributions of the minimum cloud base heights recorded by each of the Auger lidars. What is interesting about the figure is the apparent excess of minimum cloud base heights distributed between $\sim$ 8–10 km a.s.l. at Los Leones and $\sim$ 6–8 km a.s.l. at Coihueco. Clouds populated at those heights are related to the downwind turbulence of the Andes mountains [117].

Both Figure 6.8 and Figure 6.9 illustrate that the clouds observed at Auger can be influenced by the surrounding climate. Future work could perhaps investigate the relationship using data collected with our infrared cloud cameras (such as determining whether there is a pattern to cloud movement across the array). Additionally, images from the optical cameras installed at the fluorescence detector sites could be analysed. Whilst our thermal imaging cameras allow us to study clouds during the night, the optical cameras can provide us with information throughout the day.

Chapter 7

Conclusions

The major focus of this thesis has been the thermal imaging cameras at the Pierre Auger Observatory, which are used as cloud monitors. The current generation cameras are radiometric, and are sensitive to long-wave infrared radiation within the 8–14 µm waveband. Not only is atmospheric absorption and emission low at these wavelengths, but the majority of clouds are efficient thermal radiators. As a result, clouds often appear bright and stand out against the dim clear sky background. Hence, our camera systems are well-suited for night-time cloud detection at the Observatory.

Our cloud cameras are housed in a weather shield and view the sky through a protective window. They are installed on the roof of each of the main fluorescence detector (FD) buildings at Auger, where they are programmed to scan the night sky at regular intervals. The sequence of images captured throughout a given scan are transferred to Adelaide, and analysed to determine whether clouds are located within the fluorescence detector's field of view. Since clouds can adversely impact upon some of the extensive air shower measurements made using the fluorescence technique, potentially affected data are often removed from cosmic ray analyses. An important goal of my studies has been to generate "masks" that provide the cloud cover contained within each FD pixel. This information is stored in the cloud camera database, which is one of the atmospheric databases maintained by the Observatory.

Besides scene temperatures, a camera's output depends on other factors, such as the software settings and operating temperature. Compensating for the latter is important in order to analyse our images quantitatively (the camera settings can be controlled). This is because they are located outdoors, where they are exposed to a wide range of environmental temperatures. One way to compensate for these effects is to convert a camera's output into a temperature reading (an absolute measurement). Typically, temperature calibrations would be performed in a laboratory using an object that resembles a black body. However, since the cameras were already collecting data at the Observatory, this meant that the procedure had to be carried out remotely. I have calibrated our thermal imaging cameras (to an accuracy of about 2 K) using temperature measurements on clear nights made by a (calibrated) single-pixel infrared radiometer, which is also located at Auger and shares a somewhat similar spectral response with our cameras.

Another challenge has been that a camera's lens as well as the protective window also emit thermal radiation, which is detected by the sensor. In order to increase sensitivity to cloud thermal emission, it was desirable to compensate for these effects. This was achieved using flat-field images (the camera's response to a uniform source of illumination) by exposing the camera to thermal radiation from the zenith on clear nights. The correction also removes non-uniformities due to variations in the pixel-to-pixel sensitivity. These data processing routines to calibrate our cameras and improve the image quality have been derived using data collected on clear nights. It would be possible to monitor the performance of these procedures in the future by adding more clear night data to the existing routines. One important finding has been that the techniques have been reproducible for each of our cameras. Since many of the challenges presented to us have not been unique to a specific camera model (for example, thermal emission due to the optics, variations in pixel responses, the need for a temperature calibration, etc.), these practices could perhaps benefit other researchers in the field.

Clouds can often be identified through visual inspection of the thermal images, but the large amounts of data collected by our cameras means that it is necessary to automate the process. I identify clouds after the data have been processed. Key to their detection is understanding thermal emission from clear night skies. Although atmospheric absorption and emission is minimised at 8–14 µm wavelengths, the increasing path lengths of water vapour and carbon dioxide (greenhouse gases) with increasing zenith angle means that the clear night sky brightness varies. The high opacity and brightness of the atmosphere towards the horizon is one of the disadvantages of ground-based thermal detectors.

A simple method that can be used to identify clouds in our images is to subtract the clear sky emission (the background) from our measurements. The residual signal can often be attributed to clouds. I have found the best results by determining the background directly from the data when possible rather than estimating it from parameterisations of clear sky models. This usually corresponds to the lower bound of pixel values, as it often reflects the coolest (and clearest) regions of the sky. An empirical formula that can describe the change in effective sky temperature with zenith angle reasonably well is fitted to the lower bound as an estimate for the background. Some limitations of this method of cloud identification are the detection of optically thin and/or high altitude clouds, as they have a rather low contribution to the overall sky brightness, and appear only faintly warmer (or brighter) than the clear sky background. Another challenge is the detection of clouds near the horizon, as the empirical formula does not describe our camera data at larger zenith angles well (greater than about 87°).

In this dissertation, many characteristics of thermal radiation for various atmospheric conditions have been illustrated using the radiative transfer model, MODTRAN. One of the program's features is that it allows users to describe an atmosphere using height-dependent profiles. This information is readily available to us from the Global Data Assimilation System (GDAS). In this sense, the radiative transfer program can be used to simulate the clear night sky background we would expect within our cameras' band-pass. Better knowledge of the background signal could perhaps increase our sensitivity to clouds, particularly towards the horizon. Although some preliminary results have been promising, the inaccurate descriptions of GDAS surface temperatures (through comparison to local weather stations) currently lead to discrepancies between simulations and cloud camera observations. Correcting the atmospheric profiles for this effect will be the subject of future work.

The research featured in this thesis has also shown that the GDAS atmospheric temperature and humidity profiles can be used to infer whether the sky is likely to be clear or cloudy above the Observatory. This was determined by applying a technique that had originally been used to predict cloud layers from radiosonde measurements [113]. Through comparisons with the lidars operating at Auger, findings suggest that below 8 km a.s.l. about 80 % of clear skies are successfully predicted, whilst about 60 % of cloudy skies are successfully predicted. Since the GDAS data are available to members of the collaboration [21], a routine (presented in this dissertation) could be implemented as a supplementary check for the atmospheric conditions.

Appendix A

MODerate resolution atmospheric TRANsmission (MODTRAN)

MODTRAN (MODerate resolution atmospheric TRANsmission) is a computer program designed to model atmospheric propagation of electromagnetic radiation from ultraviolet through to far-infrared wavelengths (0.2 µm–100 µm). The program was developed and continues to be maintained through a collaboration between Spectral Sciences, Inc. and the U.S. Air Force. The software version that I have used in this dissertation is MODTRAN5. Supplementary reading material may be found in the MODTRAN report [118] and the user's manual [119].

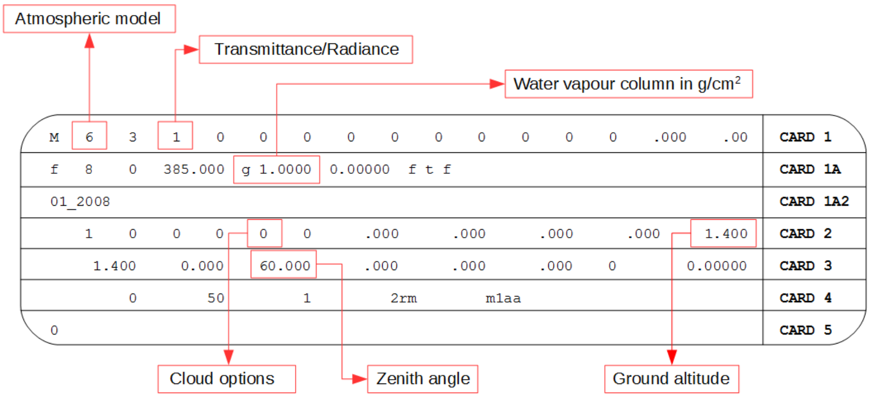

```
M    6    3    1    0    0    0    0    0    0    0    0    0    0     .000     .00   CARD 1
f    8    0  385.000 g 1.0000 0.00000  f t f                                        CARD 1A
01_2008                                                                             CARD 1A2
     1    0    0    0    0    0      .000      .000      .000      .000     1.400   CARD 2
       1.400     0.000    60.000      .000      .000      .000    0        0.00000   CARD 3
          0        50         1         2rm        m1aa                             CARD 4
0                                                                                   CARD 5
```

Figure A.1: An example of a `tape5` file used to provide inputs for MODTRAN. Each card (listed to the right) handles a separate aspect of the program and atmosphere, and some of the data fields are labelled.

The radiative transfer program calculates atmospheric radiance and transmittance at a moderate resolution (a spectral resolution as fine as $0.1\,\text{cm}^{-1}$ in MODTRAN5) along a line-of-sight specified by the user, and takes into account absorption, emission, and scattering characteristics of the atmosphere. The user defines various parameters with a single input file (named `tape5`). The file resembles a card deck, and each card (and sub-card) handles various aspects of the atmosphere and MODTRAN. Figure A.1 shows an example of the `tape5` file format, and highlights several of the parameters that I have frequently adjusted in my studies (the others were often set to default values). The main cards are:

- **CARD 1** – Main radiative transport driver.
- **CARD 1A** – Radiative transport driver (continued).
- **CARD 2** – Aerosol and cloud options.
- **CARD 3** – Line-of-sight geometry.
- **CARD 4** – Spectral range and resolution.

There are 6 predefined atmospheres that can be used in MODTRAN which describe different geographical locations and/or seasons. In addition, various types of aerosols can be selected. To illustrate the effects that various atmospheric conditions can have at long-wave infrared wavelengths, the majority of plots found in this dissertation use the 1976 U.S. Standard Atmosphere with the rural aerosol model (with 23 km visibility). One of the features in MODTRAN is that it also allows users to define an atmosphere by specifying the temperature, pressure, and humidity at different altitudes. Height-dependent atmospheric profiles can be obtained from radiosonde data, as well as from the Global Data Assimilation System (see GDAS in Section 3.6). Preliminary results using the latter to define an atmosphere in MODTRAN in order to simulate clear night sky radiation within the band-pass of our cameras may be found in Section 5.4.

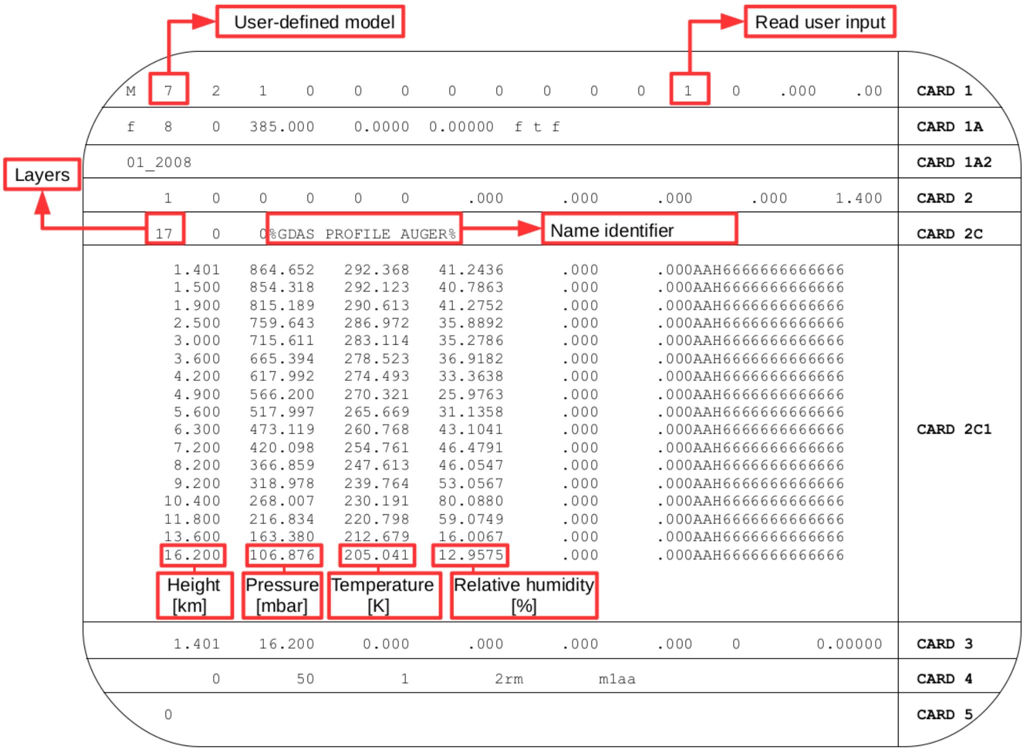

```
M    7    2    1    0    0    0    0    0    0    0    0    1    0   .000     .00
f    8    0  385.000    0.0000  0.00000  f t f
01_2008
     1    0    0    0    0    0      .000      .000      .000      .000     1.400
    17    0    0%GDAS PROFILE AUGER%
     1.401   864.652   292.368   41.2436      .000      .000AAH6666666666666
     1.500   854.318   292.123   40.7863      .000      .000AAH6666666666666
     1.900   815.189   290.613   41.2752      .000      .000AAH6666666666666
     2.500   759.643   286.972   35.8892      .000      .000AAH6666666666666
     3.000   715.611   283.114   35.2786      .000      .000AAH6666666666666
     3.600   665.394   278.523   36.9182      .000      .000AAH6666666666666
     4.200   617.992   274.493   33.3638      .000      .000AAH6666666666666
     4.900   566.200   270.321   25.9763      .000      .000AAH6666666666666
     5.600   517.997   265.669   31.1358      .000      .000AAH6666666666666
     6.300   473.119   260.768   43.1041      .000      .000AAH6666666666666
     7.200   420.098   254.761   46.4791      .000      .000AAH6666666666666
     8.200   366.859   247.613   46.0547      .000      .000AAH6666666666666
     9.200   318.978   239.764   53.0567      .000      .000AAH6666666666666
    10.400   268.007   230.191   80.0880      .000      .000AAH6666666666666
    11.800   216.834   220.798   59.0749      .000      .000AAH6666666666666
    13.600   163.380   212.679   16.0067      .000      .000AAH6666666666666
    16.200   106.876   205.041   12.9575      .000      .000AAH6666666666666
     1.401    16.200     0.000      .000      .000      .000    0       0.00000
         0        50         1       2rm      m1aa
    0
```

Figure A.2: An example of a `tape5` file with a user-defined atmosphere. Some of the data fields unique to user-defined atmospheres are labelled.

Appendix B

A list of major changes to the cloud camera software

Since the camera settings can affect measurements made with our cloud cameras, some care needs to be taken when comparing images from different epochs (times when the settings were different). The table below lists some of the major changes to the cloud camera software, and the approximate dates those changes were made.

Date	**Change(s)**
<	Auto mode. Camera settings were not saved.
2013-05-26	VSK = 3550, span = 1, and level = 65536.
2013-08-01	VSK = 3550, span = 25, and level = 45000.
2013-09-24	VSK = 3550, span = 1, and level = 65536.
2013-11-06	VSK = 3500, span = 5, and level = 40000.
2014-03-20	VSK = 3550, span = 1, and level = 65536. Values are now fixed.
2014-06-19	Temperature of the camera's sensor (T_{cam}) is now saved.
2014-07-17	Full sky mosaics went from 23 images to 19 images.
2014-10-12	Went from 384×284 to 384×288 pixel images.
	The top 4 pixels rows are saturated.
	Raw ADC counts now saved. Older images have poorer resolution.
	Shutter images are now saved, but incorrectly.

Table B.1 Continued:

2015-05-05	Shutter images are correctly saved.
	First and last images in a sequence affected by the shutter.
2016-03-08	Fixed the shutter from affecting the first/last images.
2016-10-08	The images in a given sequence had different offsets.
2016-11-01	Original cloud camera data after this period was deleted.
2017-01-10	Fix to stop the data from deleting.
2017-01-24	Fixed the different image offsets.

Table B.1: Major changes to the cloud camera software. The dates can vary slightly depending on the cloud camera. The VSK relates to the voltage bias of the blind microbolometers, the span controls the image contrast, and the level controls the image brightness (see Section 3.1.2.2).

Apart from experimentation in 2013 and at the beginning of 2014, the camera settings are now kept constant. Further information regarding the cloud cameras may be found in Section 3.1.2.

In addition to the software changes listed above, pixel pointing directions for the Coihueco camera have changed at particular times. One of the changes occurred at 2015-11-14 03:46:52 UTC due to the camera's pointing calibration file resetting. The zenith angles for each pixel were found to be more accurate after this occurred, and data collected prior to that date can be corrected by subtracting $\Delta\theta$ (listed in Table B.2) from the recorded values.

Period	$\boldsymbol{\Delta\theta}$
< 2014-10-15 23:40:32	2.2°
2014-10-15 23:40:32 to 2015-05-06 09:52:01	2.8°
2015-05-06 09:52:01 to 2015-11-14 03:46:52	5.3°
> 2015-11-14 03:46:52	0°

Table B.2: Adjustments that can be made to the zenith pointing directions for the Coihueco cloud camera depending on the time (in UTC). This is achieved by subtracting $\Delta\theta$ from the recorded values.

Appendix C

Temperature calibration results

This appendix contains some of the intermediate steps when the temperature calibrations were developed for the Auger infrared cloud cameras. Details regarding the method may be found in Section 4.5. To summarise the results of the temperature calibration:

1. Our camera's digital output (in ADC counts) not only depends on the temperature of an object in the field of view (T_{obj} in Kelvin), but also the temperature of the camera's sensor (T_{cam} in Kelvin). Some examples are shown in Figure C.1. The output can be converted into a temperature reading using the following equation:

$$\mathrm{ADC} = m\left(T_{cam}\right) \times T_{obj} + C\left(T_{cam}\right) . \tag{4.2}$$

2. $m\left(T_{cam}\right)$ can be parameterised (see Figure C.2) using:

$$m\left(T_{cam}\right) = m_2 \times T_{cam}^2 + m_1 \times T_{cam} + m_0 . \tag{4.13}$$

3. $C\left(T_{cam}\right)$ can be parameterised (see Figure C.3) using:

$$C\left(T_{cam}\right) = C_2 \times T_{cam}^2 + C_1 \times T_{cam} + C_0 . \tag{4.14}$$

4. However, the predicted temperatures (T_{obj} in Kelvin) were found to disagree with known temperatures (T_{true} in Kelvin). The temperature differences are given by:

$$T_{obj} - T_{true} = r\left(T_{cam}\right) . \tag{4.16}$$

5. $r\left(T_{cam}\right)$ can be parameterised (see Figure C.4) using:

$$r\left(T_{cam}\right) = r_3 \times T_{cam}^3 + r_2 \times T_{cam}^2 + r_1 \times T_{cam} + r_0\,. \tag{4.17}$$

6. Hence, a more accurate temperature reading can be calculated using the following:

$$T_{obj} = \frac{\mathrm{ADC} - C\left(T_{cam}\right)}{m\left(T_{cam}\right)} - r\left(T_{cam}\right)\,. \tag{4.18}$$

Table 4.3 lists the parameters for $m(T_{cam})$, Table 4.4 lists the parameters for $C(T_{cam})$, and Table 4.5 lists the parameters for $r(T_{cam})$.

Overall, the error in the temperature calibration (calculated using Equation 4.19) is about 1.9 K for the Los Leones camera, 2.1 K for the Los Morados camera, 1.8 K for the Loma Amarilla camera, and 2.5 K for the Coihueco camera. Since the camera's software settings also affect the digital output, these temperature calibrations work reasonably well for data collected after June 19th 2014, which is when we started to record the temperature of the camera's sensor, and the camera settings were fixed (refer to Appendix B).

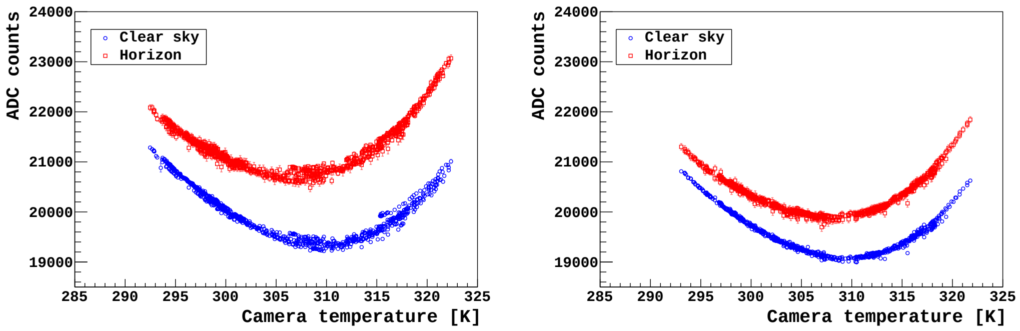

(a) Los Leones **(b)** Los Morados

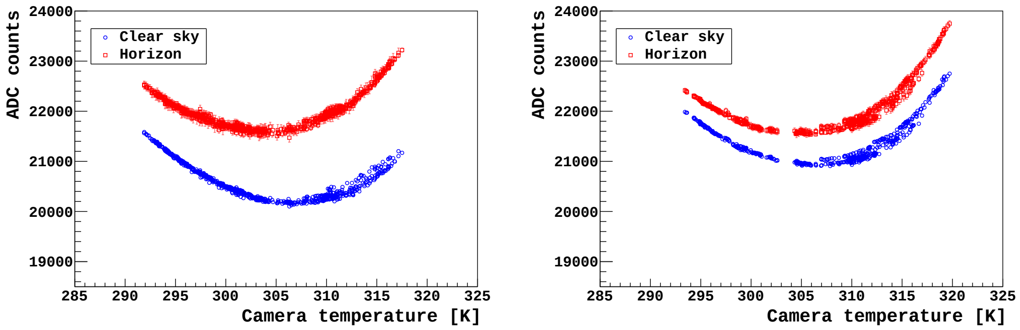

(c) Loma Amarilla **(d)** Coihueco

Figure C.1: Measurements made by the Auger infrared cloud cameras at different camera sensor temperatures (T_{cam}). Data were collected from full-sky scans performed on clear nights. The cameras viewed two types of scenes (at different effective temperatures). One was a cool scene of a clear sky near the zenith (blue circle markers). The second was a warmer scene near the horizon (red square markers). Mean pixel values for each scene are plotted, and errors derived from the standard deviation of those averaged values. The cameras' digital outputs (in ADC counts) for the two different scenes can be used to determine temperature calibrations, given that corresponding temperatures are known from a single-pixel radiometer also operating at the Pierre Auger Observatory.

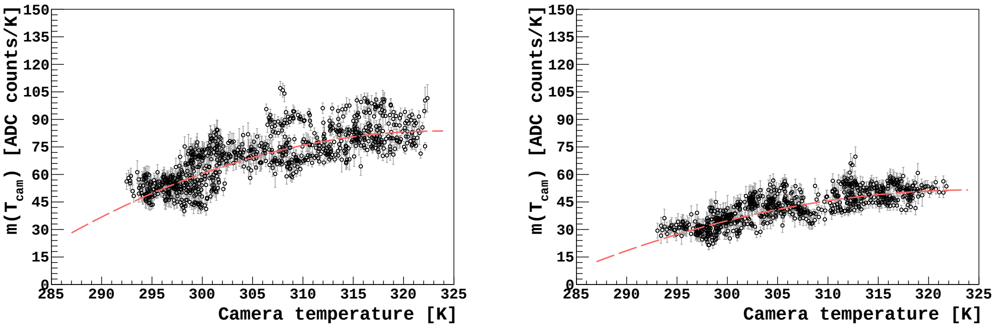

(a) Los Leones **(b)** Los Morados

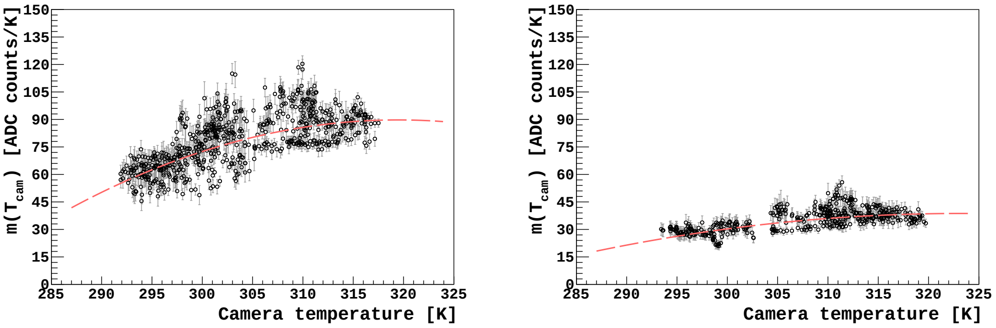

(c) Loma Amarilla **(d)** Coihueco

Figure C.2: The temperature calibration slopes $m(T_{cam})$ plotted against the temperature of the camera's sensor (T_{cam}). Results are for each of the Auger cloud cameras, and have been collected by performing two-point calibrations on clear nights using known temperatures from a single-pixel radiometer. For each camera, the dependence can be parameterised using a parabolic function (see Equation 4.13). The fits to the data are indicated by the dashed red parabolas, and a summary of the fit parameters are listed in Table 4.3.

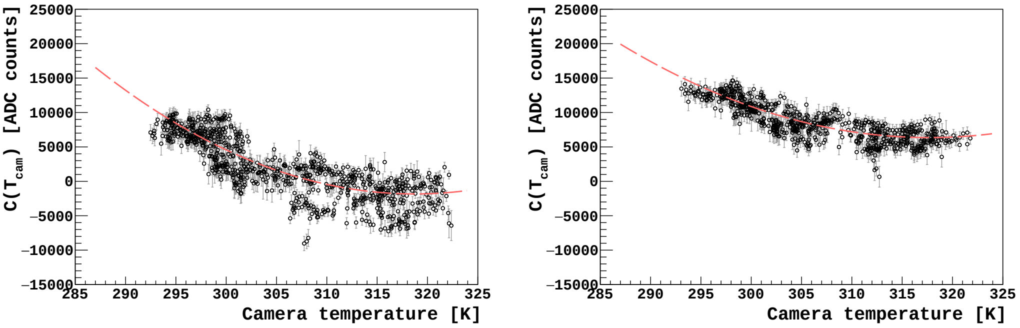

(a) Los Leones **(b)** Los Morados

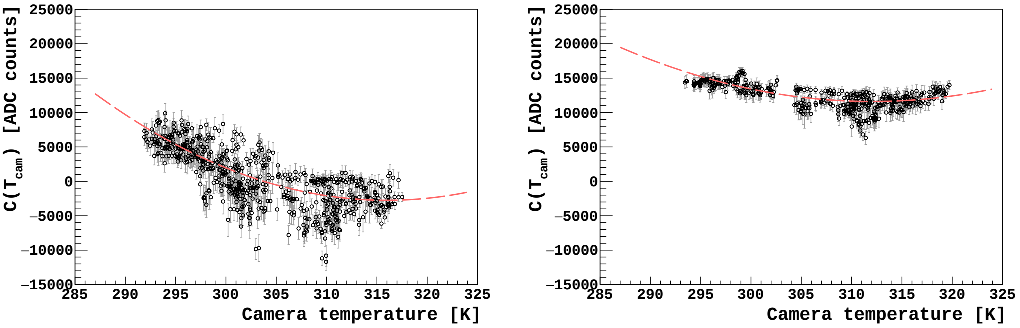

(c) Loma Amarilla **(d)** Coihueco

Figure C.3: The temperature calibration offsets $C(T_{cam})$ plotted against the temperature of the camera's sensor (T_{cam}). Results are for each of the Auger cloud cameras, and have been collected by performing two-point calibrations on clear nights using known temperatures from a single-pixel radiometer. For each camera, the dependence can be parameterised using a parabolic function (see Equation 4.14). The fits to the data are indicated by the dashed red parabolas, and a summary of the fit parameters are listed in Table 4.4.

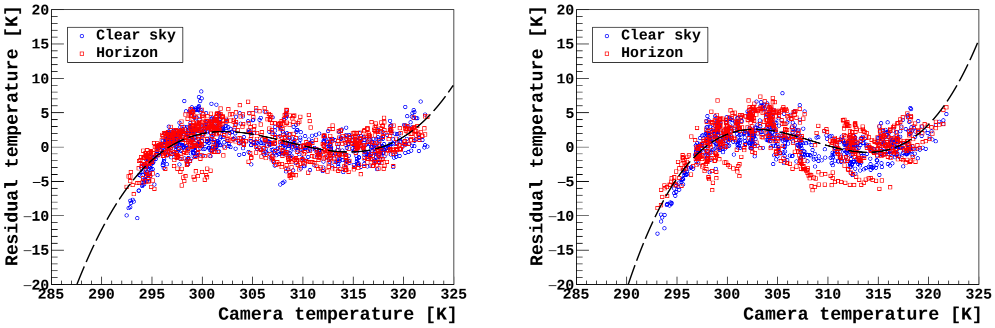

(a) Los Leones **(b)** Los Morados

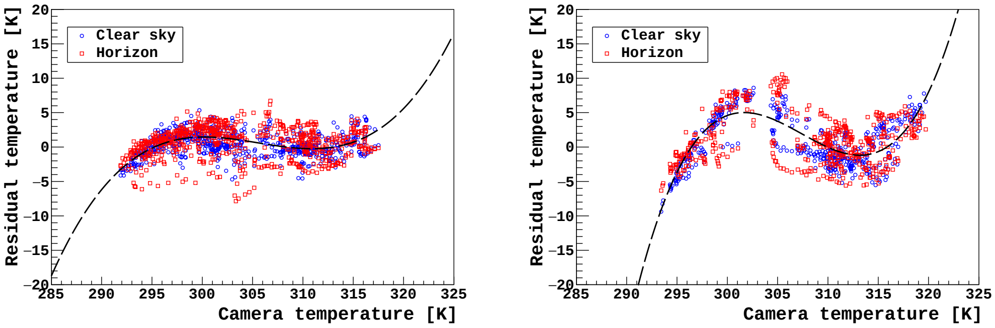

(c) Loma Amarilla **(d)** Coihueco

Figure C.4: Temperature differences plotted against the temperature of the camera's sensor (T_{cam}). The residual temperature is the calibrated temperature using camera data (a prediction based on Equation 4.2) minus the temperature measurement made by a single-pixel radiometer (the assumed truth). There are two types of temperature measurements recorded by the radiometer. One is an overhead sky temperature (and compared to overhead camera data), which is given by the blue circle markers. The second is a thermistor temperature, which can be used as a proxy for the local air temperature (and compared to camera data near the horizon), and given by the red square markers. The results show the residuals depend on T_{cam}, and suggests that modifications can be made to the temperature calibrations (see Equation 4.16). For each camera, the dependence can be parameterised using a cubic function (see Equation 4.17). The fits to the data are indicated by the dashed black curves, and a summary of the fit parameters are listed in Table 4.5.

Appendix D

A routine to check for clear/cloudy skies using GDAS atmospheric profiles

This procedure determines whether a sky is likely to be clear or cloudy based on the atmospheric temperature and humidity profiles at that time. The original technique was developed to predict cloud layers from inflections in radiosonde temperature and relative humidity measurements [113]. I substitute this data with information from the GDAS model.

Below 8 km a.s.l. $\sim 80\,\%$ of clear skies are correctly predicted. On the other hand, cloudy skies are successfully determined $\sim 60\,\%$ of the time. These results come from comparisons made to the four lidar stations operating at the Pierre Auger Observatory. Further details regarding this research may be found in Section 6.2. The method is as follows:

1. Read in height-dependent temperature $T(h)$, relative humidity $R(h)$, and dew point depression profiles. Use the approximate height resolution of the GDAS model (listed in Table 3.1).

2. Calculate the derivatives of the temperature and relative humidity vertical profiles, with respect to height. This is done using numerical differentiation. I approximate the derivatives as slopes between adjacent points, and compute the slopes using step sizes Δh. The step sizes are determined by the height resolution of the model.

3. Similarly, calculate the second derivatives of the temperature $T^{''}(h)$ and relative humidity $R^{''}(h)$ vertical profiles, with respect to height.

4. Find all regions where $R^{''}(h) \leq 0$ and $T^{''}(h) \geq 0$. Regions that satisfy these criteria represent possible cloud layers.

5. Within each predicted cloud layer, determine the minimum dew point depression.

6. Using the minimum dew point depression, along with the corresponding air temperature at that height, predict the cloud cover of each layer from the Arabey diagram (see Figure 6.4).

7. Record the cloud base height, thickness, and cloud cover of each predicted cloud layer.

8. Consider all cloud layers below 8 km a.s.l. If a layer has a cloud cover $\leq 20\,\%$ then it is considered to be clear. Similarly, when the cloud cover is greater than 20 %, but with a layer thickness ≤ 300 m, it is also said to be clear. Only when all cloud layers are clear, is the atmosphere diagnosed as a clear sky. Otherwise the sky is cloudy.

Below is a sample program that I have written in C++, which determines whether the sky is likely to be clear or cloudy.

```
//**************************************************
//A routine to test whether a sky is likely to be clear or ↩
    cloudy.
//Uses temperature and humidity data from the GDAS model at the↩
     Pierre Auger Observatory.
//80% agreement with the lidars that it is clear below 8 km a.s↩
    .l.
//60% agreement with the lidars that cloud is below 8 km a.s.l.
//Author: Patrick van Bodegom.
//Date:   24th Dec 2017.
//**************************************************

//C++ header files.
#include <iostream>
#include <fstream>

//ROOT header file.
#include "TGraph.h"

//*****************Other*functions******************

//Calculate the saturated vapour pressure (in hPa) using the ↩
    Magnus formula. Input temperature is in degrees Celsius.
double calc_svp ( double t );

//Calculate the dew point temperature (Celsius). Uses partial ↩
    pressure of water vapour (in hPa).
```

```
double calc_tdew( double e );

//Find the cloud cover (%) using the Arabey diagram. Input is temperature (t in Celsius) and dew point depression (dd in Celsius).
//Dew point depression is air temperature minus the dew point.
int arabey ( double t , double dd );

//*****************Main*function******************

int main() {

  //Read in the GDAS data.
  ifstream in_file( "gdas_file.dat" );

  //Plots for temperature, temperature first derivative, and temperature second derivative.
  TGraph g_cloud_t , g_cloud_dt , g_cloud_d2t;

  //Plots for relative humidity, relative humidity first derivative, and relative humidity second derivative.
  TGraph g_cloud_r , g_cloud_dr , g_cloud_d2r;

  //Plots of temperature, dew point depression, temperature second derivative, and relative humidity second derivative, with heights on the x-axis.
  TGraph g_t_eval , g_dd_eval;    //To estimate the cloud cover with the Arabey diagram.
  TGraph g_d2t_eval , g_d2r_eval; //To test for inflections in temperature and relative humidity.

  //Counter for the plots.
  int count = 0;

  //Vectors that contain cloud base height, cloud thickness, and the cloud cover for each predicted cloud layer.
  //All these quantities are used to check whether the sky is likely to be clear/cloudy.
  std::vector<double> vcloud_height;     //All cloud base heights predicted in the GDAS profile (m).
  std::vector<double> vcloud_thickness;  //The thickness of each cloud layer (m).
  std::vector<int>    vcloud_cover;      //Estimated cloud cover using the Arabey diagram (%).

  //GDAS data fields.
  //Height a.s.l. (m) , temperature (K) , pressure (hPa) , vapourpressure (hPa).
  double height , temperature , pressure , vapourpressure;
```

```

  //Read each line from the input file.
  while ( in_file >> height >> temperature >> pressure >> ↩
     vapourpressure ) {

    //Use the approximate height resolution of the GDAS model.
    //Surface level + 23 upper levels. Approximate heights from↩
        1976 U.S. Standard Atmosphere.
    if ( height == 1405   ||
       height == 100    ||
       height == 300    ||
       height == 500    ||
       height == 800    ||
       height == 1000   ||
       height == 1500   ||
       height == 1900   ||
       height == 2500   ||
       height == 3000   ||
       height == 3600   ||
       height == 4200   ||
       height == 4900   ||
       height == 5600   ||
       height == 6300   ||
       height == 7200   ||
       height == 8200   ||
       height == 9200   ||
       height == 10400  ||
       height == 11800  ||
       height == 13600  ||
       height == 16200  ||
       height == 20600  ||
       height == 26400
       ) {

      //Calculate the saturated vapour pressure (hPa).
      double svp = calc_svp( temperature - 273.15 );

      //Convert water vapour pressure into relative humidity ↩
         (%).
      double rhumidity = ( vapourpressure / svp ) * 100.;

      //Calculate the dew point temperature.
      double dewpoint = calc_tdew( vapourpressure );

      //Plot the temperature (in C) and the relative humidity.
      //These values are used later on to predict the cloud ↩
         layers.
      g_cloud_t.SetPoint( count , temperature - 273.15 , height↩
          / 1000. );
```

```
    g_cloud_r.SetPoint( count , rhumidity , height / 1000. );

    //Plot heights on the x-axis for these values, so they can be interpolated from later on.
    //Used to predict cloud cover from the Arabey diagram.
    g_t_eval.SetPoint( count , height / 1000. , temperature - 273.15 );
    g_dd_eval.SetPoint( count , height / 1000. , temperature - 273.15 - dewpoint );

    ++count;

    //End of the GDAS height resolution.
  }

  //End of reading the GDAS data.
}


//The method predicts cloud layers from inflections in temperature and humidity profiles.
//Calculate first derivatives, which are only used to determine the second derivatives.
//Loop through points in the graph.
for ( int point = 0; point < g_cloud_t.GetN() - 1; ++point ) {

  //Heights for this point (i) and the following (f).
  double h_i , h_f;

  //Temperature for this point and the following.
  double t_i , t_f;

  //Relative Humidity for this point and the following.
  double r_i , r_f;

  //Get temperature points from the graph.
  g_cloud_t.GetPoint( point , t_i , h_i );
  g_cloud_t.GetPoint( point + 1 , t_f , h_f );

  //Get relative humidity points from the graph.
  g_cloud_r.GetPoint( point , r_i , h_i );
  g_cloud_r.GetPoint( point + 1 , r_f , h_f );

  //First derivatives.
  double h = ( h_i + h_f ) / 2.;              //Average height between the 2 points.
  double dt = ( t_f - t_i ) / ( h_f - h_i );  //The temperature derivative.
```

```
    double dr = ( r_f - r_i ) / ( h_f - h_i );  //The relative ↩
       humidity derivative.

    //Plot.
    g_cloud_dt.SetPoint( point , dt , h );
    g_cloud_dr.SetPoint( point , dr , h );

    //End of looping over points.
  }


  //Calculate second derivatives.
  //Loop through points in the graph.
  for ( int point = 0; point < g_cloud_dt.GetN() - 1; ++point )↩
      {

    //Heights for this point (i) and the following (f).
    double h_i , h_f;

    //Derivative of temperature for this point and the ↩
       following.
    double dt_i , dt_f;

    //Derivative of relative humidity for this point and the ↩
       following.
    double dr_i , dr_f;

    //Get temperature points from the graph.
    g_cloud_dt.GetPoint( point , dt_i , h_i );
    g_cloud_dt.GetPoint( point + 1 , dt_f , h_f );

    //Get relative humidity points from the graph.
    g_cloud_dr.GetPoint( point , dr_i , h_i );
    g_cloud_dr.GetPoint( point + 1 , dr_f , h_f );

    //Second derivative.
    double h = ( h_i + h_f ) / 2.;                    //Average ↩
       height between the 2 points.
    double d2t = ( dt_f - dt_i ) / ( h_f - h_i );  //The ↩
       temperature second derivative.
    double d2r = ( dr_f - dr_i ) / ( h_f - h_i );  //The ↩
       relative humidity second derivative.

    //Plot height on the x-axis for values, so they can be ↩
       interpolated in the next step.
    g_d2t_eval.SetPoint( point , h , d2t );
    g_d2r_eval.SetPoint( point , h , d2r );

    //End of looping over points.
```

```
}


//To test for cloud layers, check where the criteria ( d2r <= 0 and d2t >= 0 ) are met.
//Check if this is the first measurement of the cloud layer (i.e. the base).
bool isbase = false;

//Check if this is the last measurement of the cloud layer (i.e. the top).
bool istop = false;

//The heights for the cloud base and cloud top (m).
double cloud_base = -99999 , cloud_top = -99999;

//Find the minimum dew point depression in a cloud layer, and the corresponding temperature.
//This is used for the Arabey diagram to predict cloud coverage.
double ddmin = 99999 , tmin = 99999;

//Interpolate values within the graphs up to a height of 14 km a.s.l.
for ( int point = 0; point <= 1000; ++point ) {

  //The height (km).
  double h = 0 + point * ( 14 - 0 ) / 1000.;

  //Temperature (C).
  double t = g_t_eval.Eval( h );

  //Dew point depression (C).
  double dd = g_dd_eval.Eval( h );

  //Temperature second derivative (C/km^2).
  double d2t = g_d2t_eval.Eval( h );

  //Relative humidity second derivative (%/km^2).
  double d2r = g_d2r_eval.Eval( h );

  //Initially assume no cloud at this height.
  bool iscloud = false;

  //Check if the measurement is above ground level (1401 m).
  //Not interested in the data below these heights.
  if ( h >= 1401. / 1000. ) {

    //Check criteria for a possible cloud layer.
```

```
     if ( d2r <= 0 && d2t >= 0 ) {

  //If conditions are met, then cloud was predicted at this ↩
     height.
  iscloud = true;

  //Search for the minimum dew point temperature and ↩
     corresponding temperature in the cloud layer.
  if ( dd < ddmin ) {
    ddmin = dd;
    tmin  = t;
    //End of searching for minimum values.
  }

  //End of cloud conditions criteria.
     }

     //Check if this is the top of the cloud layer.
     //Test the conditions for cloud are no longer met, and a ↩
        cloud base exists, and this is the last time cloud ↩
        was found in the layer.
     else if ( !( d2r <= 0 && d2t >= 0 ) && isbase == true && ↩
        istop == false ) {

  //Record the cloud top height.
  cloud_top = h;

  //The cloud top has been found.
  istop = true;

  //End of testing for the top of the cloud layer.
     }

     //End of ground level check.
   }

   //Check if cloud was found, and this was the first time it ↩
      was found in the layer.
   if ( iscloud == true && isbase == false ) {

     //Record the cloud base height.
     cloud_base = h;

     //The cloud base has been found.
     isbase = true;

     //End of testing for the cloud base.
   }

```

```
    //When the cloud has both a base and top, record the value.↩
        Then reset the booleans to test for further cloud ↩
        layers.
    if ( isbase == true && istop == true ) {

      //The thickness of the cloud layer (m).
      double cloud_thickness = ( cloud_top - cloud_base );

      //Predicted cloud coverage (%) using the Arabey diagram.
      int cloud_cover = arabey ( tmin , ddmin );

      //Record the cloud base height (m) , cloud layer ↩
          thickness (m) and the cloud cover (%).
      vcloud_height.push_back( cloud_base * 1000. );
      vcloud_thickness.push_back( cloud_thickness * 1000. );
      vcloud_cover.push_back( cloud_cover );

      //Reset booleans.
      isbase = false;
      istop  = false;

      //Reset the minimum dew point temperature and ↩
          corresponding temperature.
      ddmin = 99999;
      tmin  = 99999;

      //End of checking the cloud base and top exist.
    }

    //End of looping over heights.
  }


  //Use information from all the predicted cloud layers to ↩
      determine whether the sky is likely to be clear or cloudy↩
      .
  //Total number of predicted GDAS cloud layers.
  int nlayers = vcloud_height.size();

  //Initially assume that it is not clear.
  bool gdas_clear = false;

  //A cloud layer is considered to be clear when the cloud ↩
      cover is less than 20%, or if the cover is greater than ↩
      20%, but the thickness is <= 300 m.
  int NP1 = 0; //Number of layers with cloud cover <= 20%. (↩
      Clear instance 1).
  int NP2 = 0; //Number of layers with cloud cover > 20% and a ↩
      thickness <= 300 metres. (Clear instance 2).
```

```

  //Only consider the heights below 8 km a.s.l., as clouds are almost always predicted above those heights.
  int nlayers8km = 0; //Number of layers below 8 km a.s.l.

  //Loop over all cloud layers.
  for ( int g = 0; g < nlayers; ++g ) {

    double height_g = vcloud_height[g];     //Cloud base height (m).
    double thick_g  = vcloud_thickness[g]; //Cloud thickness (m).
    int    cov_g    = vcloud_cover[g];      //Coverage (%).

    //Only consider layers below 8 km a.s.l.
    if ( height_g < 8000. ) {

      //Clear instance 1.
      if ( cov_g <= 20 ) ++NP1;

      //Clear instance 2.
      if ( cov_g > 20 && thick_g <= 300. ) ++NP2;

      //Increment layers.
      ++nlayers8km;

      //End of layers below 8 km.
    }

    //End of looping over GDAS cloud layers.
  }

  //Change the boolean if it was found to be clear up to 8 km a.s.l.
  //That is, all predicted cloud layers were clear.
  if ( NP1 + NP2 == nlayers8km ) gdas_clear = true;

  //Display whether the sky is likely to be clear or overcast.
  if ( gdas_clear == true )  std::cout << "There is an 80% chance the sky is clear" << std::endl;
  if ( gdas_clear == false ) std::cout << "There is a 60% chance the sky is cloudy" << std::endl;

  return 0;

  //End of the main function.
}


```

```
//****************Other*functions*****************

//Calculate the saturated vapour pressure (in hPa) using the Magnus formula. Input temperature is in degrees Celsius.
double calc_svp ( double t ) {

  //Saturated vapour pressure (in hPa).
  double e_s = 0.;

  //The Magnus formula.
  //When temperature is >= 0.
  if ( t >= 0 ) e_s = 6.1070 * exp( ( 17.15 * t ) / ( 234.9 + t ) );

  //When temperature is < 0.
  if ( t < 0 ) e_s = 6.1064 * exp( ( 21.88 * t ) / ( 265.5 + t ) );

  //Return the saturated vapour pressure (svp).
  return e_s;

  //End of calculating the saturated vapour pressure.
}



//Calculate the dew point temperature (Celsius). Uses the partial pressure of water vapour (in hPa).
double calc_tdew ( double e ) {

  //Plot the saturated vapour pressure (hPa) against temperature (Celsius) using the Magnus formula.
  TGraph g_mag;

  //Loop over points, covering the range -60 - 60 degrees Celsius.
  for ( int point = 0; point <= 600; ++point ) {

    double x = -60. + point * ( 60 - -60. ) / 600.; //Temperature.
    double y = calc_svp( x );                       //Saturated vapour pressure.

    g_mag.SetPoint( point , y , x );

  }
```

```

  //Find the dew point temperature. This is when e = es.
  double tdew = g_mag.Eval( e );

  //Return the dew point.
  return tdew;

  //End of calculating the dew point temperature.
}



//Find the cloud cover (%) using the Arabey diagram. Input is temperature (t in Celsius) and dew point depression (dd in Celsius).
//Dew point depression is air temperature minus the dew point.
int arabey ( double t , double dd ) {

  //The cloud coverage (%).
  int coverage = -99999;

  //Four zones in the diagram that give the cloud cover (%).
  //Zone #1 : 80 - 100 %.
  //Zone #2 : 60 - 80 %.
  //Zone #3 : 20 - 60 %.
  //Zone #4 :  0 - 20 %.

  //The boundary lines for each zone ( 1-2 , 2-3 and 3-4 ).
  //The boundary between zone 1 and 2.
  TGraph g_zone12;
  g_zone12.SetPoint( 0 , -80. , 8.00 );  //Extrapolate the line back to -80 C.
  g_zone12.SetPoint( 1 , -40. , 4.00 );
  g_zone12.SetPoint( 2 , -10. , 1.00 );
  g_zone12.SetPoint( 3 ,   0. , 0.81 );
  g_zone12.SetPoint( 4 ,  40. , 0.81 );

  //The boundary between zone 2 and 3.
  TGraph g_zone23;
  g_zone23.SetPoint( 0 , -80. , 10.98 ); //Extrapolate the line back to -80 C.
  g_zone23.SetPoint( 1 , -40. , 6.10 );
  g_zone23.SetPoint( 2 , -10. , 2.44 );
  g_zone23.SetPoint( 3 ,   0. , 2.00 );
  g_zone23.SetPoint( 4 ,  40. , 2.00 );

  //The boundary between zone 3 and 4.
  TGraph g_zone34;
  g_zone34.SetPoint( 0 , -80. , 14.26 );  //Extrapolate the
```

```
     line back to -80 C.
  g_zone34.SetPoint( 1 , -40. , 8.26 );
  g_zone34.SetPoint( 2 , -10. , 3.76 );
  g_zone34.SetPoint( 3 ,   0. , 2.89 );
  g_zone34.SetPoint( 4 ,  40. , 2.89 );

  //Interpolate boundaries at x-value (the temperature) to find↩
      the y-value (dew point depression).
  double y12 = g_zone12.Eval( t );
  double y23 = g_zone23.Eval( t );
  double y34 = g_zone34.Eval( t );

  //The upper limit on cloud cover (%).
  if ( dd < y12 ) coverage = 100;
  if ( dd >= y12 && dd < y23 ) coverage = 80;
  if ( dd >= y23 && dd < y34 ) coverage = 60;
  if ( dd >= y34 ) coverage = 20;

  //Return the cloud cover.
  return coverage;

  //End of using the Arabey diagram.
}
```

References

[1] The Pierre Auger Collaboration. The Pierre Auger Cosmic Ray Observatory. *ArXiv e-prints*, February 2015.

[2] V. F. Hess. Über beobachtungen der durchdringenden strahlung bei sieben freiballonfahrten. *Physik. Zeitschr.*, 13:1084–1091, November 1912.

[3] J. Blümer, R. Engel, and J. R. Hörandel. Cosmic rays from the knee to the highest energies. *Progress in Particle and Nuclear Physics*, 63:293–338, October 2009.

[4] C. Patrignani, P. Richardson, et al. Review of particle physics. *Chin. Phys. C*, 40:100001, 2016.

[5] K. Greisen. End to the cosmic-ray spectrum? *Physical Review Letters*, 16:748–750, April 1966.

[6] G. T. Zatsepin and V. A. Kuz'min. Upper limit of the spectrum of cosmic rays. *Soviet Journal of Experimental and Theoretical Physics Letters*, 4:78, August 1966.

[7] A. A. Penzias and R. W. Wilson. A measurement of excess antenna temperature at 4080 Mc/s. *Astrophys. J.*, 142:419–421, July 1965.

[8] J. W. Cronin. Cosmic rays: the most energetic particles in the universe. *Reviews of Modern Physics Supplement*, 71:165, March 1999.

[9] The Pierre Auger Collaboration. Observation of a large-scale anisotropy in the arrival directions of cosmic rays above 8×10^{18} eV. *Science*, 357:1266–1270, September 2017.

[10] A. Haungs, H. Rebel, and M. Roth. Energy spectrum and mass composition of high-energy cosmic rays. *Reports on Progress in Physics*, 66:1145–1206, July 2003.

[11] The Pierre Auger Collaboration. The Pierre Auger Observatory Upgrade - Preliminary Design Report. *ArXiv e-prints*, April 2016.

[12] The Pierre Auger Collaboration. Observation of the suppression of the flux of cosmic rays above 4×10^{19} eV. *Physical Review Letters*, 101:061101, August 2008.

[13] The Pierre Auger Collaboration. Upper limit on the cosmic-ray photon fraction at EeV energies from the Pierre Auger Observatory. *Astroparticle Physics*, 31:399–406, July 2009.

[14] The Pierre Auger Collaboration. Improved limit to the diffuse flux of ultrahigh energy neutrinos from the Pierre Auger Observatory. *Physical Review D*, 91(9):092008, May 2015.

[15] The Pierre Auger Collaboration. Depth of maximum of air-shower profiles at the Pierre Auger Observatory. II. Composition implications. *Phys. Rev.*, D90(12):122006, 2014.

[16] I. Allekotte, A. F. Barbosa, P. Bauleo, et al. The surface detector system of the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 586:409–420, March 2008.

[17] T. Stanev. *High Energy Cosmic Rays*. Springer Praxis Books. Springer Berlin Heidelberg, 2010.

[18] M. Ave, M. Bohacova, B. Buonomo, et al. Measurement of the pressure dependence of air fluorescence emission induced by electrons. *Astroparticle physics*, 28:41–57, September 2007.

[19] J. Abraham, P. Abreu, M. Aglietta, et al. The fluorescence detector of the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 620:227–251, August 2010.

[20] F. Arqueros, J. R. Hörandel, and B. Keilhauer. Air fluorescence relevant for cosmic-ray detection—Summary of the 5th fluorescence workshop, El Escorial 2007. *Nuclear Instruments and Methods in Physics Research A*, 597:1–22, November 2008.

[21] The Pierre Auger Collaboration. Description of atmospheric conditions at the Pierre Auger Observatory using the Global Data Assimilation System (GDAS). *Astroparticle Physics*, 35:591–607, April 2012.

[22] M. G. Winnick. *Cloud cameras at the Pierre Auger Observatory*. PhD thesis, University of Adelaide, 2010.

[23] J. Chirinos for the Pierre Auger Collaboration. Remote sensing of clouds using satellites, lidars, CLF/XLF and IR cameras at the Pierre Auger Observatory. *EPJ Web of Conferences*, 89:03012, December 2015.

[24] J. Abraham, P. Abreu, M. Aglietta, et al. A study of the effect of molecular and aerosol conditions in the atmosphere on air fluorescence measurements at the Pierre Auger Observatory. *Astroparticle Physics*, 33:108–129, March 2010.

[25] V. Ramanathan, R. D. Cess, E. F. Harrison, et al. Cloud-radiative forcing and climate: Results from the Earth Radiation Budget Experiment. *Science*, 243:57–63, January 1989.

[26] D. L. Hartmann. *Global Physical Climatology*. International geophysics series. Academic Press, 1994.

[27] R. Sloan, J. H. Shaw, and D. Williams. Infrared emission spectrum of the atmosphere. *Journal of the Optical Society of America*, 45:455, June 1955.

[28] J. R. Allen. Measurements of cloud emissivity in the 8–13 μm waveband. *Journal of Applied Meteorology*, 10:260–265, April 1971.

[29] M. Vollmer and K.P. Möllmann. *Infrared Thermal Imaging: Fundamentals, Research and Applications*. Wiley, 2010.

[30] K. N. Liou. *An Introduction to Atmospheric Radiation*. International Geophysics. Elsevier Science, 2002.

[31] W. S. Benedict, H. H. Claassen, and J. H. Shaw. Absorption spectrum of water vapor between 4.5 and 13 microns. *Journal of Research of the National Bureau of Standards*, 49(2), 1952.

[32] J. L. Monteith and M. H. Unsworth. *Principles of Environmental Physics: Plants, Animals, and the Atmosphere*. Elsevier Science, 2013.

[33] G. B. Rybicki and A. P. Lightman. *Radiative Processes in Astrophysics*. Physics textbook. Wiley, 2008.

[34] O. A. Alduchov and R. E. Eskridge. Improved Magnus form approximation of saturation vapor pressure. *Journal of Applied Meteorology*, 35:601–609, April 1996.

[35] H. R. Pruppacher and J. D. Klett. *Microphysics of Clouds and Precipitation*. Atmospheric and Oceanographic Sciences Library. Springer Netherlands, 2010.

[36] R. A. Houze. *Cloud Dynamics*. International Geophysics. Elsevier Science, 1994.

[37] A. Otárola, T. Travouillon, M. Schöck, et al. Thirty Meter Telescope site testing X: Precipitable water vapor. *Publications of the Astronomical Society of the Pacific*, 122:470, April 2010.

[38] A. Maghrabi and H. M. Al Dajani. Estimation of precipitable water vapour using vapour pressure and air temperature in an arid region in central Saudi Arabia. *Journal of the Association of Arab Universities for Basic and Applied Sciences*, 14(1):1–8, 2013.

[39] A. Otárola, R. Querel, and F. Kerber. Precipitable water vapor: Considerations on the water vapor scale height, dry bias of the radiosonde humidity sensors, and spatial and temporal variability of the humidity field. *ArXiv e-prints*, March 2011.

[40] S. Jin and O. F. Luo. Variability and climatology of PWV from global 13-year GPS observations. *IEEE Transactions on Geoscience and Remote Sensing*, 47:1918–1924, July 2009.

[41] M. D. King. Radiative properties of clouds. *International Geophysics*, 54:123–149, 1993.

[42] E. Bauer. The scattering of infrared radiation from clouds. *Applied Optics*, 3(2):197–202, February 1964.

[43] G. E. Hunt. Radiative properties of terrestial clouds at visible and infra-red thermal window wavelengths. *Quarterly Journal of the Royal Meteorological Society*, 99:346–369, April 1973.

[44] Z. Sun and K. P. Shine. Studies of the radiative properties of ice and mixed-phase clouds. *Quarterly Journal of the Royal Meteorological Society*, 120:111–137, January 1994.

[45] D. Riordan, R. Clay, A. Maghrabi, et al. Cloud base temperature measurements using a simple longwave infrared cloud detection system. *Journal of Geophysical Research (Atmospheres)*, 110:D03207, February 2005.

[46] The Pierre Auger Collaboration. Identifying clouds over the Pierre Auger Observatory using infrared satellite data. *Astroparticle Physics*, 50:92–101, December 2013.

[47] Xenics Infrared Solutions: Gobi-384 infrared camera. http://www.xenics.com/sites/default/files/leaflets/xb-015_04_gobi-384_25um_scientific_highres.pdf. (Last accessed May 8th, 2018).

[48] J. L. Tissot, C. Trouilleau, B. Fieque, et al. Uncooled microbolometer detector: recent developments at ULIS. *Opto-Electronics Review*, 14:25–32, March 2006.

[49] C. Minassian, J. L. Tissot, M. Vilain, et al. Uncooled amorphous silicon TEC-less 1/4 VGA IRFPA with 25 μm pixel-pitch for high volume applications. In *Infrared Technology and Applications XXXIV*, volume 6940 of *Proceedings of SPIE*, page 69401Z, April 2008.

[50] H.VIR® windows. http://www.hvir.com/hvir_infos.php. (Last accessed May 8th, 2018).

[51] Xenics Infrared Solutions. Private communication.

[52] R. W. Clay, N. R. Wild, D. J. Bird, et al. A cloud monitoring system for remote sites. *Publications of the Astronomical Society of Australia*, 15:332–335, November 1998.

[53] A. H. Maghrabi. *Ground based measurements of atmospheric infrared radiation from clear and cloudy skies.* PhD thesis, University of Adelaide, 2007.

[54] KUBE Electronics AG: Single PIR lenses. http://www.kube.ch/downloads/pdf/kube_single_fresnel_lenses.pdf. (Last accessed May 8th, 2018).

[55] PerkinElmer Optoelectronics: TPS 534 datasheet. http://www.alldatasheet.com/datasheet-pdf/pdf/14942/PERKINELMER/TPS534.html. (Last accessed May 8th, 2018).

[56] A. Maghrabi and R. Clay. Precipitable water vapour estimation on the basis of sky temperatures measured by a single-pixel IR detector and screen temperatures under clear skies. *Meteorological Applications*, 17:279–286, December 2009.

[57] National Oceanic and Atmospheric Administration (NOAA): GOES-8 through GOES-15 Imager spectral response functions. `http://www.ospo.noaa.gov/Operations/GOES/goes-imager-srfs.html`. (Last accessed May 8th, 2018).

[58] National Oceanic and Atmospheric Administration (NOAA): Change of channels on Imagers beginning with GOES-12. `http://www.ospo.noaa.gov/Operations/GOES/calibration/change-channels.html`. (Last accessed May 8th, 2018).

[59] G. P. Ellrod. Advances in the detection and analysis of fog at night using GOES multispectral infrared imagery. *Weather and Forecasting*, 10:606–619, September 1995.

[60] J. Chirinos for the Pierre Auger Collaboration. Ground-truthing a satellite-based night-time cloud identification technique at the Pierre Auger Observatory. *European Physical Journal Plus*, 127:93, August 2012.

[61] S. Y. Benzvi, R. Cester, M. Chiosso, et al. The lidar system of the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 574:171–184, April 2007.

[62] A. Tonachini. *Atmospheric calibration of the fluorescence detectors with the lidar system of the Pierre Auger Observatory*. PhD thesis, University of Torino, 2007.

[63] B. Fick, M. Malek, J. A. J Matthews, et al. The Central Laser Facility at the Pierre Auger Observatory. *Journal of Instrumentation*, 1(11):P11003, 2006.

[64] The Pierre Auger Collaboration. Techniques for measuring aerosol attenuation using the Central Laser Facility at the Pierre Auger Observatory. *Journal of Instrumentation*, 8:P04009, April 2013.

[65] L. Valore. *Atmospheric Aerosol Determination using Vertical Laser Tracks in the Pierre Auger Experiment*. PhD thesis, Università degli Studi di Napoli "FedericoII", 2007.

[66] B. Wilczyńska, D. Góra, P. Homola, et al. Variation of atmospheric depth profile on different time scales. *Astroparticle Physics*, 25:106–117, March 2006.

[67] The Pierre Auger Collaboration. Impact of atmospheric effects on the energy reconstruction of air showers observed by the surface detectors of the Pierre Auger Observatory. *Journal of Instrumentation*, 12(02):P02006, 2017.

[68] The AIRFLY Collaboration. Temperature and humidity dependence of air fluorescence yield measured by AIRFLY. *Nuclear Instruments and Methods in Physics Research A*, 597:50–54, November 2008.

[69] M. Unger, B. R. Dawson, R. Engel, et al. Reconstruction of longitudinal profiles of ultra-high energy cosmic ray showers from fluorescence and Cherenkov light measurements. *Nuclear Instruments and Methods in Physics Research A*, 588:433–441, April 2008.

[70] J. C. Owens. Optical refractive index of air: dependence on pressure, temperature and composition. *Applied Optics*, 6(1):51–59, 1967.

[71] K. P. Birch and M. J. Downs. An updated Edlén equation for the refractive index of air. *Metrologia*, 30:155–162, January 1993.

[72] A. Bucholtz. Rayleigh-scattering calculations for the terrestrial atmosphere. *Applied Optics*, 34(15):2765–2773, May 1995.

[73] The Pierre Auger Collaboration. Atmospheric effects on extensive air showers observed with the surface detector of the Pierre Auger Observatory. *Astroparticle Physics*, 32:89–99, September 2009.

[74] National Oceanic and Atmospheric Administration (NOAA): Global Data Assimilation System (GDAS). `https://www.ncdc.noaa.gov/data-access/model-data/model-datasets/global-data-assimilation-system-gdas`. (Last accessed May 9th, 2018).

[75] D. Epperlein, B. Keilhauer, J. Blümer, et al. Investigation of applying a global atmospheric model to the southern site of the Pierre Auger Observatory. Private communication (GAP-2010-074, 2010. Pierre Auger Collaboration internal technical note).

[76] M. Will for the Pierre Auger Collaboration. Global atmospheric models for cosmic ray detectors. *ArXiv e-prints*, February 2014.

[77] I. Durre, R. S. Vose, and D. B. Wuertz. Overview of the Integrated Global Radiosonde Archive. *Journal of Climate*, 19:53–68, January 2006.

[78] NOAA Air Resources Laboratory (ARL): Global Data Assimilation System (GDAS1) Archive Information. https://www.ready.noaa.gov/gdas1.php. (Last accessed May 9th, 2018).

[79] M. Hori, T. Aoki, T. Tanikawa, et al. In-situ measured spectral directional emissivity of snow and ice in the 8-14 μm atmospheric window. *Remote Sensing of Environment*, 100:486–502, February 2006.

[80] P. Mullinger and B. Jenkins. *Industrial and Process Furnaces: Principles, Design and Operation*. Butterworth-Heinemann/IChemE series. Butterworth-Heinemann, 2008.

[81] G. Bieszczad, T. Orżanowski, et al. Method of detectors offset correction in thermovision camera with uncooled microbolometric focal plane array. In *Electro-Optical and Infrared Systems: Technology and Applications VI*, volume 7481 of *Proceedings of SPIE*, page 74810O, September 2009.

[82] A. Durand, C. Minassian, et al. Uncooled amorphous silicon TEC-less 1/4 VGA IRFPA with 25μm pixel-pitch for high volume applications. In *Infrared Technology and Applications XXXV*, volume 7298 of *Proceedings of SPIE*, page 72982B, May 2009.

[83] P. W. Nugent, J. A. Shaw, and N. J. Pust. Correcting for focal-plane-array temperature dependence in microbolometer infrared cameras lacking thermal stabilization. *Optical Engineering*, 52(6):061304, June 2013.

[84] J. L. Marshall and D. L. DePoy. Flattening scientific CCD imaging data with a dome flat field system. *ArXiv Astrophysics e-prints*, October 2005.

[85] P. Kask, K. Palo, C. Hinnah, and T. Pommerencke. Flat field correction for high-throughput imaging of fluorescent samples. *Journal of Microscopy*, 263(3):328–340, September 2016.

[86] A. V. Arecchi, R. J. Koshel, and T. Messadi. *Field Guide to Illumination*. Field Guide Series. SPIE, 2007.

[87] N. Menn. *Practical Optics*. Electronics & Electrical. Elsevier Academic Press, 2004.

[88] J. M. Mooney and F. D. Shepherd. Characterizing IR FPA nonuniformity and IR camera spatial noise. *Infrared Physics and Technology*, 37:595–606, August 1996.

[89] A. Daniels. *Field Guide to Infrared Systems*. SPIE field guides. Society of Photo Optical, 2007.

[90] FLIR Systems Inc. The Ultimate Infrared Handbook for R&D Professionals. https://www.flirmedia.com/MMC/THG/Brochures/T559243/T559243_EN.pdf. (Last accessed May 9th, 2018).

[91] A. F. Milton, F. R. Barone, and M. R. Kruer. Influence of nonuniformity on infrared focal plane array performance. *Optical Engineering*, 24:855, August 1985.

[92] Y. Tendero, S. Landeau, and J. Gilles. Non-uniformity correction of infrared images by midway equalization. *Image Processing On Line*, 2:134–146, 2012.

[93] J. A. Sobrino and J. Cuenca. Angular variation of thermal infrared emissivity for some natural surfaces from experimental measurements. *Applied Optics*, 38:3931–3936, June 1999.

[94] M. Martin and P. Berdahl. Summary of results from the spectral and angular sky radiation measurement program. *Solar Energy*, 33:241–252, January 1984.

[95] A. J. Prata. A new long-wave formula for estimating downward clear-sky radiation at the surface. *Quarterly Journal of the Royal Meteorological Society*, 122:1127–1151, July 1996.

[96] X. Berger, D. Buriot, and F. Garnier. About the equivalent radiative temperature for clear skies. *Solar Energy*, 32:725–733, January 1984.

[97] W. C. Swinbank. Long-wave radiation from clear skies. *Quarterly Journal of the Royal Meteorological Society*, 89:339–348, July 1963.

[98] T. M. Crawford and C. E. Duchon. An improved parameterization for estimating effective atmospheric emissivity for use in calculating daytime downwelling longwave radiation. *Journal of Applied Meteorology*, 38:474–480, April 1999.

[99] M. H. Unsworth and J. L. Monteith. Long-wave radiation at the ground I. Angular distribution of incoming radiation. *Quarterly Journal of the Royal Meteorological Society*, 101:13–24, January 1975.

[100] C. N. Awanou. Clear sky emissivity as a function of the zenith direction. *Renewable Energy*, 13(2):227–248, 1998.

[101] S. Smith and R. Toumi. Measuring cloud cover and brightness temperature with a ground-based thermal infrared camera. *Journal of Applied Meteorology and Climatology*, 47:683, 2008.

[102] J. A. Shaw, P. W. Nugent, N. J. Pust, et al. Radiometric cloud imaging with an uncooled microbolometer thermal infrared camera. *Optics Express*, 13(15):5807–5817, Jul 2005.

[103] J. A. Shaw and P. W. Nugent. Physics principles in radiometric infrared imaging of clouds in the atmosphere. *European Journal of Physics*, 34(6):S111, 2013.

[104] M. K. Daniel and P. M. Chadwick. Atmospheric monitoring with an infrared radiometer. In *EPJ Web of Conferences*, volume 89, page 03005. EDP Sciences, 2015.

[105] W.X. Zhang, D.R. LÜ, and Y.L. Chang. A feasibility study of cloud base height remote sensing by simulating ground-based thermal infrared brightness temperature measurements. *Chinese Journal of Geophysics*, 50(2):339–350, March 2007.

[106] L. Liu, X.J. Sun, X.C. Liu, et al. Comparison of cloud base height derived from a ground-based infrared cloud measurement and two ceilometers. *Advances in Meteorology*, 2015, 2015.

[107] S. W. Seemann, E. E. Borbas, R. O. Knuteson, et al. Development of a global infrared land surface emissivity database for application to clear sky sounding retrievals from multispectral satellite radiance measurements. *Journal of Applied Meteorology and Climatology*, 47:108, 2008.

[108] S. K. Cox. Cirrus Clouds and the Climate. *Journal of Atmospheric Sciences*, 28:1513–1514, November 1971.

[109] C. M. R. Platt. Lidar and radiometric observations of cirrus clouds. *Journal of Atmospheric Sciences*, 30:1191–1204, September 1973.

[110] J. P. Craven, R. E. Jewell, and H. E. Brooks. Comparison between observed convective cloud-base heights and lifting condensation level for two different lifted parcels. *Weather and Forecasting*, 17:885–890, August 2002.

[111] J. Zhang, Z. Li, H. Chen, and M. Cribb. Validation of a radiosonde-based cloud layer detection method against a ground-based remote

sensing method at multiple ARM sites. *Journal of Geophysical Research (Atmospheres)*, 118:846–858, January 2013.

[112] J. Wang and W. B. Rossow. Determination of cloud vertical structure from upper-air observations. *Journal of Applied Meteorology*, 34:2243–2258, October 1995.

[113] I. V. Chernykh and R. E. Eskridge. Determination of cloud amount and level from radiosonde soundings. *Journal of Applied Meteorology*, 35:1362–1369, August 1996.

[114] E. N. Arabey. Radiosonde data as means for revealing cloud layers (Russian). *Meteor. Gidrol.*, 6:32–37, 1975.

[115] J. Day, B. Dawson, and R. Clay. Estimation of cloud base height above the Pierre Auger Observatory using GDAS. Private communication (GAP-2017-055, 2017. Pierre Auger Collaboration internal technical note).

[116] W. R. Cotton and R. A. Anthes. *Storm and Cloud Dynamics*. International Geophysics. Elsevier Science, 1992.

[117] R. Mussa. Private communication.

[118] L. W. Abreu and G. P. Anderson. The MODTRAN 2/3 report and LOWTRAN 7 model. *Contract*, 19628(91-C):0132.

[119] A. Berk, G. P. Anderson, P. K. Acharya, and E. P. Shettle. MODTRAN 5.2.0.0 user's manual. July 2008.