Overcoming Challenges of Sky Background Light in Large-Scale Astronomical Mosaics

March 2016

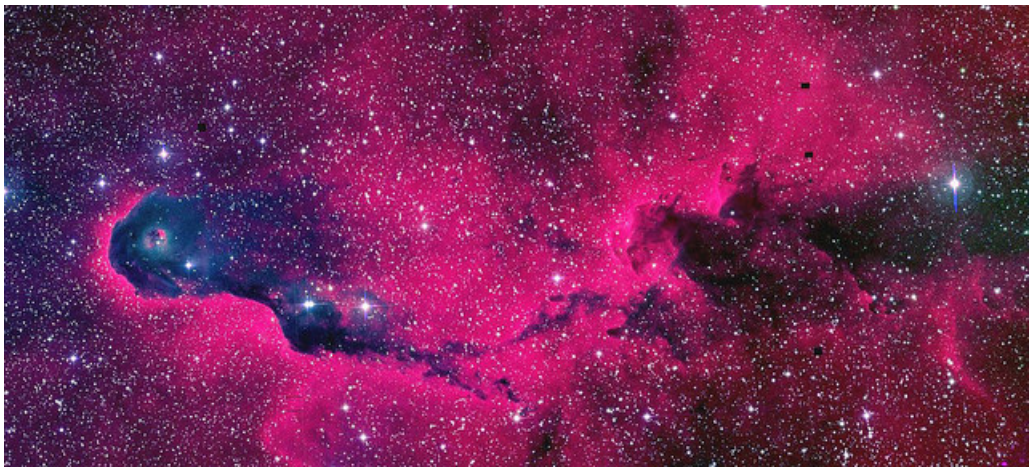

Charles William Sidney Roe

Submitted to the University of Hertfordshire in partial fulfilment of the requirements of the degree of MSc by Research.

Physics, Astronomy and Mathematics
University of Hertfordshire
United Kingdom

Overcoming Challenges of Sky Background Light in Large-Scale Astronomical Mosaics

March 2016

Charles Roe

Abstract

Large-scale, high-resolution, photometrically calibrated images are key for many astrophysical problems. The INT Photometric Hα Survey has imaged the entire northern Galactic Plane in r, i and Hα filters. However, these images suffer from a number of common imaging problems, including, most critically, large-scale gradients due to scattered moonlight. The objective of this work is to produce an automated method for cleaning this data so that it can be used to produce large-scale and reliable Hα mosaics for scientific use.

We created dark-time templates to account for airglow, fringing, and other sources of dark-time counts in the images and then used a Markov Chain Monte Carlo method to fit a linear, 2-dimensional model to the scattered moonlight. Bright stars in the images are censored from the fitted images so they do not influence the fit. Other types of model were explored, as well as a method that employed Fourier transforms to clean the data, but without fruition. The method to fit the model to the moonlight background was originally tested in the i-band, before moving onto the the r-band, subtracting scaled Hα images to remove nebulosity. An empirical scaling factor was then used to translate the model fit from the r-band to the Hα band, necessary because of varying atmospheric conditions.

Finally, the cleaned data were shifted onto a common zero point before mosaicking into large scale images. The result is a strong groundwork for cleaning astronomical images by accounting for the various components to sky background but preserving features of interest. The results of this process applied to images that cover supernova remnant Simeis 147 show a substantial improvement over uncleaned imaging data. We also illustrate the versatility of this process by applying it, unprepared, to other regions in the Galactic Plane.

Acknowledgements

First and foremost, I'd like to thank my two supervisors, who have both taken on the role of principle supervisor at various points throughout the project. Huge thanks to Dr Nick Wright and Prof. Janet Drew for their attention, time and patience throughout.

Thanks also to Dr Jim Geach, for his advice and time for the work on Fourier analysis, as well as the code he provided. Along with Jim, thanks also to all those who gave suggestions to, and comments on, the method at the various presentations.

Thanks Dr Geert Barentsen for his guidance in the first half of the project on the inner workings of IPHAS and its data releases, as well as for all his help throughout.

Big thanks to my family for their support throughout, particularly financially, to allow me to embark on this endeavour in the first place.

Finally thanks to Mena, for reminding me to stop working in the overly-enthusiastic times, and to start working in the rest.

Contents

List of Figures

1 Introduction

This work focusses on the cleaning of imaging data from wide field surveys, particularly narrow-band Hα imaging data that reveals the complexity of the interstellar medium and the processes at work within it. I therefore begin my introduction with the interstellar medium (section 1.1) and move on to the linked phenomena of young stars (1.2) and supernova remnants (1.3), before going on to briefly introducing the area of sky we have chosen to test our work. Then I will introduce the observational data itself and discuss its features.

1.1 The Interstellar Medium (ISM)

The term Interstellar Medium (ISM) encompasses all of the material in a galaxy between the stars. It is made up almost entirely of gas in various forms, while the rest of it is dust.

The ISM is intricately connected to star formation, stellar evolution, and the final death of stars. The ISM covers a wide range of physical scales, from the sub-pc scales within molecular clouds and compact HII regions, up to the tens of kpc of entire galaxies.

The most widely used model to describe the composition of the ISM splits it into 4 phases (Osterbrock, 1984), though it is often described as a 3-phase model, with the WNM and WIM making up one 'warm phase';

- Cold gas in the cold neutral medium (CNM) making up HI clouds. The CNM makes up around only 3% of the ISM by volume but contains most of its mass, and is in approximate pressure equilibrium with its surroundings (Field et al., 1969).

- Warm gas in the warm neutral and warm ionized medium (WNM/WIM) makes up roughly half of all the ISM material by volume. The WNM is typically found at the edges of HII regions and molecular clouds, while the WIM is the main component of HII regions, like giant molecular clouds. The WIM (temperatures of $\sim$ 8000K) is the primary source of Hα emission, and therefore of most interest in this work.

- Hot gas in the hot ionized medium (HIM) which fills most of the other 50% of volume of the ISM (McKee and Ostriker, 1977). The HIM is made up of overlapping supernova bubbles.

This replaced the previously accepted two phase model in Field et al. (1969), which described the ISM as being made up of only the HIM and CNM.

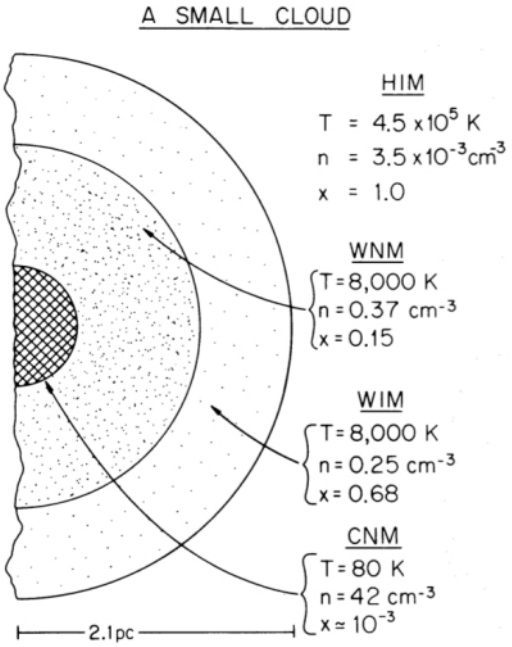

Figure 1: A diagram from McKee and Ostriker (1977) displaying the locations of the 4 main phases of the ISM in a small cloud. Typical values of hydrogen density are given as n, while the ratio of the number of ionized particles to neutral particles is x.

Mass exchange is possible between the 3 phases, as cold clouds can be heated to become diffuse gas, which can further be swept up by SN blast waves and compressed, forming high-pressure, ionized regions at the shell of supernova remnants. McKee and Ostriker (1977) proposed that the hot ionized medium is heated by supernova explosions, occurring in the dying moments of massive star evolution, that move through the dense ISM cloud with the shock front of the supernova remnants. In this way, through supernovae as well as stellar winds flowing from stars throughout their life-cycles, vast amounts of energy are pumped into the ISM, sparking turbulence and ionization of gas, as well as star formation. The ionised gas produced by these interactions with the ISM, allow us to observe and study it in detail.

While many different models of the composition of the ISM exist, another of note is the five phase model (Osterbrock, 1984). This model adds to the already set CNM, WNM,WIM and HIM with the fifth thermal phase of H_2 molecular clouds. These cold ($\sim$15K) clouds are sites of star formation that contain around a third of the total mass of the ISM, while making up only

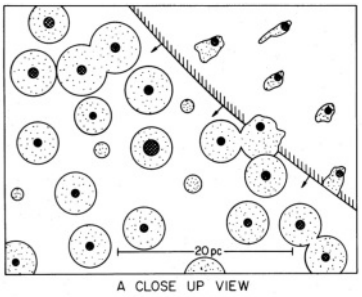

Figure 2: The sweeping up of cold clouds by a SN blastwave (the HIM) propagating from the top right, from McKee and Ostriker (1977). The blast wave sweeps up clouds and distorts their warm gas (dotted regions, the WIM/WNM) around the cold cores (dark, hatched dots, the CNM).

$\sim$0.05% of its volume. As well as molecular clouds, dust must also be included as it is a strong source of extinction and consists of all the non-gas particles within the ISM.[1]

The cold, dense clouds that make up the CNM phase are found embedded in the HIM and form the smallest phase in terms of volume in the ISM with a filling factor of $\simeq 0.025$. The typical structure of a cold cloud in McKee and Ostriker (1977)'s 3-phase model can be seen in figure 1, and shows a cold, central cloud surrounded by warm neutral and warn ionized gas from the WNM and WIM phase. This extended cloud, itself a part of the ISM, is surrounded by the HIM which makes up the majority of the volume with a filling factor of ~ 0.7.

The pressure in the ISM is maintained by supernovae (SN) in the HIM. Where supernova remnants (SNR) expand and overlap before they fade away, the ISM pressure will rise until, at the overlapping point, the pressure in the SNR is the same as the pressure in the ISM. Using this assumption, McKee and Ostriker (1977) was able to correctly predict the pressure of the ISM from the evolution of SNR in the 3-phase ISM model. As the expanding SN blast-wave propagates through the ISM, sweeping up material, it 'evaporates' the cold clouds of gas in the CNM, ionizing them and elevating their material to the HIM (Figure 2).

[1]Ohio State University course notes, P. I-25: http://www.astronomy.ohio-state.edu/ pogge/Ast871/Notes/Intro.pdf

1.2 Star formation

As the ISM cools, the cold molecular clouds collapse to form new stars of various masses. These stars fuse the abundant hydrogen into more complex elements which, in their dying moments, they expel into the surrounding ISM, enriching it. These interactions with the local ISM mean that when future star formation occurs and a protoplanetary disk forms and collapses to form planets, those planets will contain heavier elements than just hydrogen or helium, like those we see on Earth.

Pre-main sequence stars are split into 2 types; T Tauri and Herbig Ae/Be stars. T Tauri stars are pre-main sequence stars that become visible when they are able to generate sufficient energy and have dispersed enough material to be optically detected (Bertout and Bouvier, 1989). T Tauri stars are divided into two classes: Classical T Tauri stars (CTTS), which are strong accretors, and Weak-lined T Tauri stars (WTTS), which are not. About half of all T Tauri stars still have a disk of matter surrounding them, and are in the process of contracting onto the main sequence in the Hertzsprung-Russel (H-R) diagram.

Herbig Ae/Be stars are the more massive counterparts to T Tauri stars and range in mass from around 3 to $10M_{\odot}$. These intermediate mass (Waters, 2006) HAeBe stars are pre-main sequence objects since they have not yet started burning hydrogen.

High mass pre-main sequence objects (masses greater than $8M_{\odot}$) are not observed often, since they have short lived pre-main sequence life-cycles, which are over while the star is still embedded and obscured from us in the molecular cloud it was formed in. T Tauri and HAeBe stars are short-lived phases and therefore examples of them are rare, which hinders our attempts to understand their properties and their evolution. Wide area surveys, particularly those using the Hα filter, are useful for identifying CTTS due to their excess Hα emission that arises from accretion.

1.3 Supernovae and Supernova Remnants

1.3.1 The origin of supernovae

Supernovae are amongst the most luminous and energetic events in the Universe. They come in a number of varieties and originate from two major cosmic events: the thermal runaway of a white dwarf star (resulting in Type Ia supernovae) or the core collapse of a massive (> 8 $M_{\odot}$) star (resulting in Type Ib, Ic or Type II supernovae).

High mass stars (masses greater than $8M_{\odot}$) have main sequence lifetimes that are considerably shorter than low-mass stars, but are massive enough

that their core temperatures can eventually reach the high temperatures needed to fuse helium, followed by carbon, then neon, and so on, until the core is primed for the final extremely rapid burning of silicon to nickel, with layers of the remains of previous burning phases encircling it. Once nuclear burning ceases, the star is no longer in hydro-static equilibrium and it collapses at around a quarter of the speed of light, causing the outer layers to rebound off the collapsed core in a brilliant supernova explosion. These explosions return a considerable amount of material to the ISM, material that has been enriched within the progenitor star, thus enriching and replenishing the ISM. The material ejected in a supernova forms a hypersonic shock wave which continually expands with the cooling material. As the material expands and cools over time, it becomes visible at progressively longer wavelengths, allowing us to view its various properties. The remaining 'bubble' of material is a 'supernova remnant' (Weiler and Sramek, 1988).

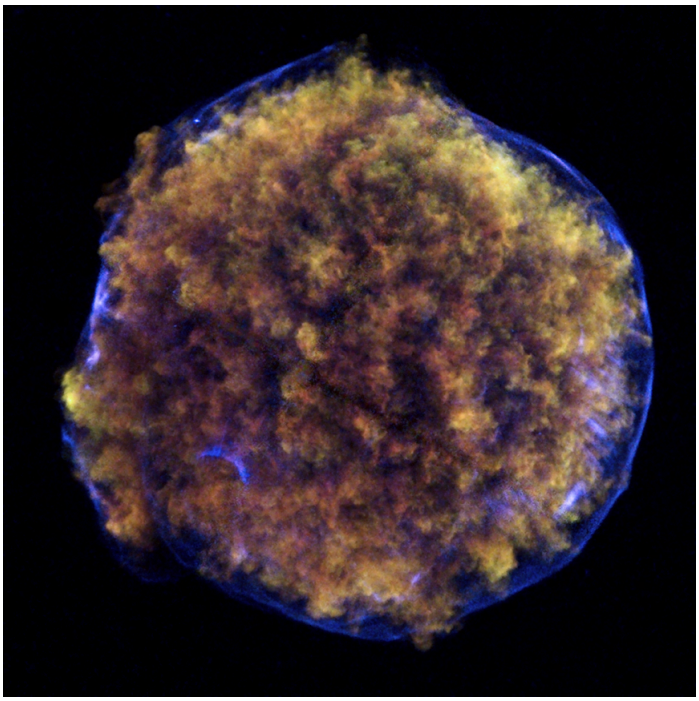

*Figure 3: A Chandra X-ray image of Tycho's supernova remnant displaying the key role X-ray astrophysics plays in studying the early stages of a SNR. Image taken from: chandra.harvard.edu/photo/*2011/*tycho*2/.

Type Ia supernova originate from a star with a mass of less than $8M_{\odot}$. Once such stars have finished burning hydrogen they will collapse, before moving on to the red giant branch (RGB) where the star burns hydrogen in a shell surrounding a degenerate helium core. Following this, it will burn

helium in a 'helium flash' before, once the star has heated up sufficiently, it begins helium burning in its core, at which point it is found on the horizontal branch (HB). When core helium burning has finished and the star is left with a carbon-oxygen core the star contracts and heats up. It is unable to reach sufficient temperatures to fuse carbon in its core, but instead will undergo an unstable process of helium burning in shells surrounding the core. During this time the star is observed on the asymptotic giant branch (AGB). During the AGB phase, the dying star sheds most of its mass in its outer layers, which are then ionized by the exposed hot core, briefly becoming visible and prominent as a planetary nebula (PN), before dispersing and the star eventually cooling as a white dwarf. White dwarfs cannot have masses higher than 1.4 $M_{\odot}$(the Chandrasekhar limit), where their densities would be high enough to overcome electron degeneracy and they would collapse into a neutron star.

Type Ia supernovae occur in a binary system where one star (the primary) is a white dwarf and the other (the secondary) is a mass-transferring 'donor' star. Mass from the donor star is accreted into a disk around the primary, making the donor star unstable, resulting in nova explosions (smaller than supernovae). The accreted matter in the disk falls steadily on to the primary, increasing its mass. In some cases, this mass increase pushes the mass of the white dwarf over the Chandrasekhar limit, resulting in a supernova explosion (Robinson, 1976). Because these type Ia supernovae are often viewed as occurring when the primary star reaches the Chandrasekhar limit, they are treated as standard candles.

1.3.2 The evolution of supernova remnants

In general, a supernova explosion, regardless of its type, follows the same set of phases (Weiler and Sramek, 1988):

- As soon as a supernova goes off, the ejecta is far denser and hotter than the circumstellar medium (CSM) around it. In fact, core collapse supernovae will have already cleared much of their environment through stellar winds and ejecta. At this point, the 'blast-wave' of ejecta propagates through the CSM at a near constant velocity causing a shock-wave as it goes. This is called the 'free expansion phase'. (Reynolds, 2008; Smartt, 2009)

- At the point at which the pressure of the shocked CSM exceeds the thermal pressure of the ejecta, a reverse shock begins to propagate backwards, moving back towards the supernova. As this reverse shock-wave moves through the ejecta, it slows and heats it. Once this reverse

shock reaches the centre and all the ejecta is very hot, the free expansion phase is over. These early, hot stages where the gas is being shock heated are traced by strong X-ray emission (Reynolds, 2008, Figure 3).

- As the expansion of the supernova remnant stalls, it is picked up by the fast radiative expansion of the hot gas interior, which leaves a cool, thin gas shell on the outside of the blast wave sweeping up material. This stage has strong optical emission, through the recombination of ionized hydrogen and oxygen atoms. This new phase is the snow plough phase, with typical shock speeds of 150km/s.

- Beyond the snow plough phase, a supernova remnant fades away as it gradually slows and its pressure declines. When the shock wave slows to the effective sound speed, it becomes a sound wave in its last state before it fades slowly away into the ISM. This phase is best observed through emission from the neutral hydrogen atoms seen in the radio.

1.3.3 Imaging and studying supernova remnants

SNRs are typically studied in the X-ray, radio and optical, with each playing an important role as the wavelengths of the peak brightness of a SNR varies over time, with its evolution. Early on in the lifetime of a SNR, while it is still a hot, dense ball of freely expanding gas, they are best viewed through high-energy X-rays. In fact, throughout the majority of a SNRs evolution, its very hot phases ($\sim$ 10000K) are most visible in the X-ray, the cool shell of swept-up ISM material in front of the shock front is most visible in the radio, while the hot, shock-ionized ISM material forming filamentary structure is most visible in the optical and, more specifically, Hα.

Wide field surveys allow us to identify new supernova remnants so we can better study their properties e.g. Sabin et al. (2013). They allow us to look at both the large scale structure of SNRs, how they have evolved, their morphology etc, as well as the very small scales of their structure, to see how they are interacting with the ISM. A current feature of interest in the area of SNRs in the literature is self-sealing shells. Self-sealing shells, or 'blowouts' (Figure 4), are described in Pittard (2013) as occurring on the surfaces of SNRs when the dense, cool shell at the blastwave confining the hot ejecta inside ruptures and the hot material is able to rapidly leak out into the ISM. This escaping ejecta sweeps up cool ISM material of its own at its outer edge and self-seals, forming a 'blister' at the surface of the SNR. Pittard (2013) suggests that ruptures can only occur while the SNR is expanding fast enough that the decelerating shell can become unstable. This initial fast expansion should affect the observed velocity structure of the shell while the

Figure 4: An example blow-out candidate as seen in the supernova remnant Simeis 147. IPHAS Hα data.

blow-out exists. There may also be links between the size of the blow-out and the size of the overall SNR shell.

Other SNR structures of interest in the literature include secondary shells (Lozinskaya, 1992). These shells exist separately from the primary blast-wave associated with an SNR and appear much fainter, found in different cases to exist both inside and outside SNRs. These secondary blast-waves could be the remains of a second supernova from the host, or potentially from a binary system. They may also be associated with the primary SNR, possibly the remains of the return shock that propagates back through the ejecta in the free-expansion phase. For these secondary shocks to be better understood, they must first be found around current SNR in high-resolution imagery.

1.3.4 Simeis 147

Producing clean, high-resolution images is the goal of this work, and Simeis 147 is chosen as the ideal candidate to test the method because of its location on the sky, its large area and highly evolved morphology, its faintness (its structure is sensitive to changes of even a couple of counts in an image) and because of the relatively large amount of lower-quality data in its area.

Simeis 147 (figure 5) - otherwise known as G180.0-01.7, Shajn 147, Sharpless 2-240 or S147 - is a late-stage supernova remnant towards the galactic anti-centre and on the galactic mid plane. Discovered in 1952 and at a distance of 0.8-1.6 kpc (Gvaramadze, 2006), S147 is a large, faint shell-type SNR on the near side of the Perseus arm that consists of numerous filaments embedded in large-scale diffuse emission. The energy released in the explosion resulting in S147 is believed to have been $\sim 11 \times 10^{50} n_0$ ergs (Silk and Wallerstein, 1973), where n_0 is the number density of the pre-shock ISM, which is typically $\sim$1 particle/cm^3.

*Figure 5: An amateur image of S147 covering ∼6×4.5 degrees. Image taken from: Rogelio Bernal Andreo, http : //apod.nasa.gov/apod/ap*121009*.html.*

Anderson et al. (1996) discovered a pulsar in S147, called PSR J0538+2817, 40 arcmin west of the centre of the SNR with an age upper limit of 6×10^5 years, consistent with the age of the remnant at 8×10^4 to 2×10^5 years. A second, runaway B0.5V star within S147 that is consistent with having been at the same position as the central pulsar approximately 30 kyr ago was found by Kramer et al. (2003), possibly part of an old binary system between it and the pulsar. The pulsar has a dispersion measured distance

of 1.8 kpc, consistent again with that of S147 at 0.8-1.6 kpc. There is an offset between the pulsar's observed characteristic age (6×10^5 years) and its kinematic age (3×10^4 years, Dinçel et al., 2015) suggested to be due to there having been two supernovae (Gvaramadze, 2006), one that gave rise to the pulsar and the other responsible for the visible SNR. This would imply that the SNR has expanded into a region of space already partly cleared by the first SN.

There is some indication that the northern half of S147 expands faster than the southern half (figure 2 of Lozinskaya, 1976), with a sharp boundary still evident in the south while being less regular in the north. The east and west edges of S147 show signs of blow-outs (see figure 5). The expansion velocity of 80-120km/s implies that the SNR has entered the final, momentum-conserving, snow plough phase of its evolution. The same conclusion can be reached from the good positional agreement between several optical and radio filaments (Sofue et al., 1980; Fuerst and Reich, 1986) and from the non-detection of X-ray emission from the SNR's shell (Sauvageot et al., 1990) .

There is seen to be high velocity gas behind S147 in Silk and Wallerstein (1973), who also shows a sequence between the SNRs Vela XYZ, the Cygnus loop and S147 in increasing size, age, expansion velocity and radio surface brightness. High electron densities in the gas in one sightline suggest a shocked interstellar cloud that has not been completely recombined (Phillips and Gondhalekar, 1983). Less depletion in the high-velocity interstellar gas than in the low-velocity gas (Phillips et al., 1981) suggests the sputtering of grains: where the high energy of the SN blast-wave has broken apart dust.

SNR become more and more unstable as they age. To properly understand the physical processes that go on inside SNRs and how SNRs affect their surroundings, they can be studied using hydrodynamic simulations that can be compared with the structure of observed SNRs.

1.4 Hα imaging surveys

Hα surveys are of a high interest in astrophysics due to the abundance of Hα light which is emitted by many stars and ionized emission nebulae (both planetary nebulae and HII regions), allowing for a better understanding of these objects. Hα is created when an electron in a hydrogen atom falls from its third to its second energy level, and has a wavelength of 656nm, which falls in the red part of the electromagnetic spectrum.

Emission nebulae are of particular interest in the eyes of this project. These regions are ionized by one or more nearby hot stars that emit a signif-

icant fraction of high-energy photons capable of ionizing hydrogen (Draine, 2011). The emission of Hα comes from the excitation of an electron in the hydrogen atom up to the third energy level in the Bohr model of the atom, before is cascades back down the energy levels, releasing photons of light. The transition from the 3rd to the 2nd energy level is responsible for Hα light, releasing a photon at its specific wavelength (656nm). The energy required to excite an electron to the 3rd energy level (hν $\sim$12eV) is around the same as that required to ionise hydrogen (hν $\geq$13.6eV). Therefore, around half of the Hα emission comes from the ionization, and the recombination of hydrogen. During recombination, the electron typically falls from a higher initial energy level and, when it cascades down from the 3rd to the 2nd levels, releases Hα .

Most of the short-lived phases of star formation and stellar evolution, as well as the phases of the ISM, can be investigated using Hα surveys since they invariably lead to enhanced Hα emission. Feedback in molecular clouds coming from various processes can also be investigated to indirectly explore stellar evolution features, such as stellar winds, the removal of circumstellar accretion disks, and the effects of supernovae on the local environment. Currently, we are lacking large samples of objects in all the phases discussed here. Hα surveys allow us to identify large numbers of objects which will help us understand how stars evolve through these phases and which can then be followed up with spectroscopy.

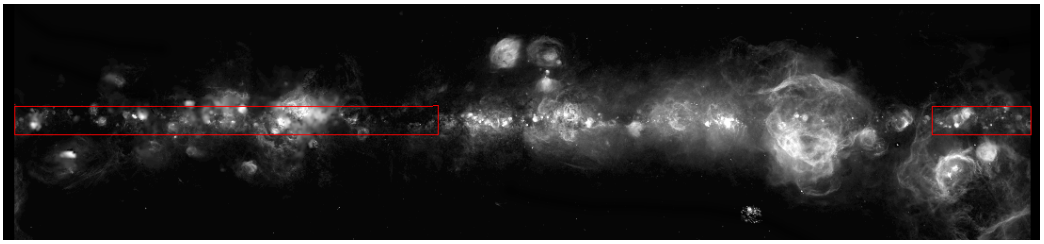

Figure 6: The Finkbeiner Composite Hα map of the Galactic plane with the area covered by IPHAS marked out by the solid red line. This image covers a galactic latitude range of $-40^\circ < b < 40^\circ$ and the entire Galactic Plane with the galactic centre in the middle. Image taken from www.skymaps.info.

1.5 IPHAS

The main focus of this report is on techniques to better prepare IPHAS Hα data for scientific exploitation.

The INT (Isaac Newton Telescope, La Palma) Photometric Hα Survey of the Northern Galactic Plane (IPHAS, Drew et al., 2005) is a 1800deg^2 survey

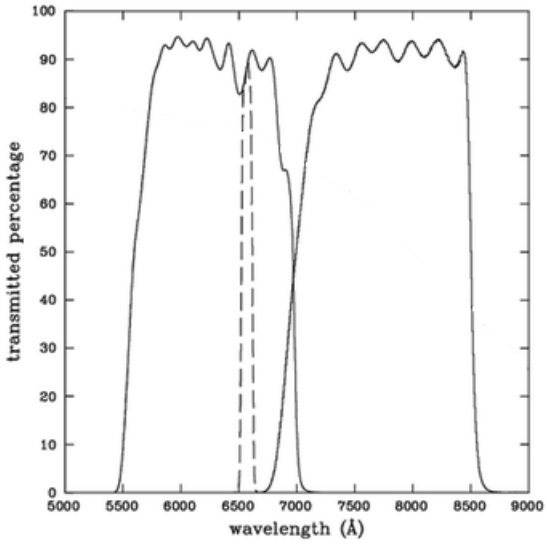

Figure 7: Transmission profiles for the 3 filters used by IPHAS. The solid lines represent the Sloan r (left) and i (right) bands, whilst Hα is plotted as the dashed line. Figure taken from Drew et al. 2005.

started in 2003, imaging in the galactic longitude range of $l = 30° - 215°$ with a latitude of $-5° < b < 5°$ (figure 6) using the Sloan r and i broad-band, and the Hα narrow-band filters (Figure 7). It has a median seeing of 1.1" and an r-band 5 σ depth of 21.2 mag. IPHAS was designed to identify objects in short phases of stellar evolution, i.e. supernova remnants, planetary nebula, T-Tauri stars, and Wolf-Rayet stars (all objects that can be efficiently identified from their excess Hα emission), as well as resolve the overall nebulosity of the ISM.

IPHAS makes use of the Wide Field Camera (WFC) on the 2.5-m INT

Figure 8: The Isaac Newton Telescope at its current site in La Palma, photo courtesy of the IAC.

(figure 8). The exposure time for images in the Hα filter are 120sec, but 30sec and 10sec for r and i respectively. The WFC comprises of 4 anti-reflection-coated, thinned 4096 x 2048 CCDs in an L-shape arrangement, with a total field of view of 0.29deg^2 (Figure 9). Each pixel has a width of 13.5μm which corresponds to an on-sky pixel dimension of 0.333 x 0.333 arcsec^2.

1.5.1 Observing strategy

Since there are not insignificant gaps between the 4 CCDs (seen in Figure 9), each pointing has a corresponding offset pointing which is displaced by 5' west and south in order to fill in the gaps between CCDs. The 2 fields that make up the footprint are called the 'on' and 'off' pointings and both together make a 'field pair' or a 'set of exposures' (Figure 10), consisting of 8 CCD frames. Each field pair is observed in all 3 filters within 10 minutes to minimize any variability between on- and off-frames and between different filters.

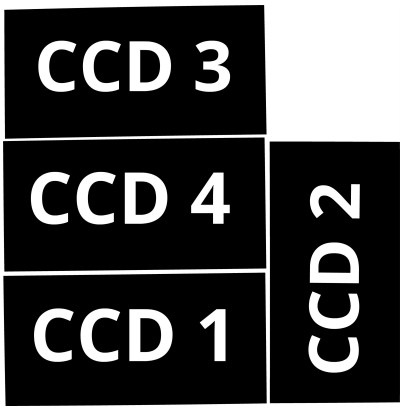

Figure 9: The setup and relative positions of the 4 CCDs making up the WFC. North is up, East is Left.

1.5.2 Data processing

IPHAS observational data are first sent to the Cambridge Astronomical Survey Unit (CASU) for initial pipeline processing. The pipeline was originally prepared for the INT Wide Field imaging survey (WFS, McMahon et al., 2001), a 200 deg^2 extragalactic survey that ran between 1998 and 2003.

IPHAS uses the same telescope and camera as the WFS and so can benefit from the same pipeline. In short, this works in 2 major steps, described in more detail in Irwin and Lewis (2001). The first is the image processing, which completes bias subtraction, initial flat fielding to take off internal gain

correction, and de-fringing of the i-band images. The second step is for object detection, which works to estimate the background counts and improve the ability to detect faint objects. A detection is when 4 or more neighbouring pixels have more counts than the number of background counts by at least 1.25σ.

1.5.3 Data releases

IPHAS has had 2 major data releases to date detailed in González-Solares et al. (2008) and Barentsen et al. (2014). The initial data release (IDR) contained 200 million unique objects found in 1600deg^2 of the footprint, whilst the second data release, in 2014, contained photometry for 219 million sources from 92% of the IPHAS footprint. The data included in this second data release (DR2) came from observations carried out from the start of the survey in 2003 up until 2012. The data contained in DR2 is not only uniformally calibrated, but contains further quality control measures than in the pipeline or the IDR.

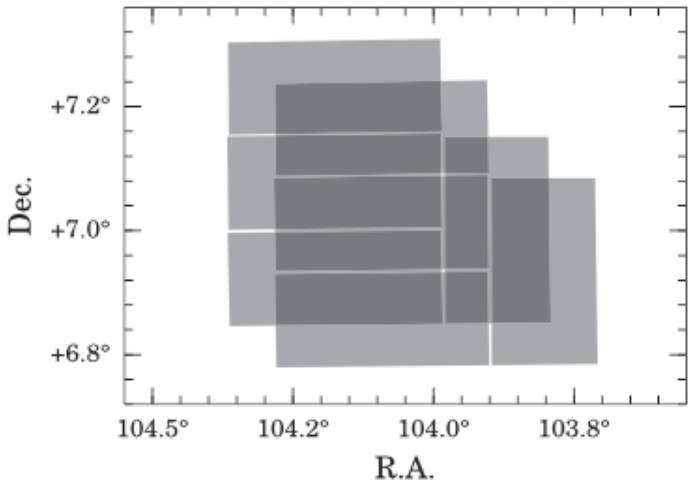

Figure 10: A field pair in the IPHAS survey, displaying the overlap of the 'on' and 'off' observations. This image comes from Barentsen et al. (2014)

More recent observations have taken place, up to the present day, as a means to replace 'bad' data. The 'bad' data is defined in Section 1.6. Some of the images within the data set known to be of particularly poor quality have been re-observed since DR2. This data, though not published, is available for use in this work and, where appropriate, has been used as replacement for particularly problematic fields. The final survey database will contain 15270

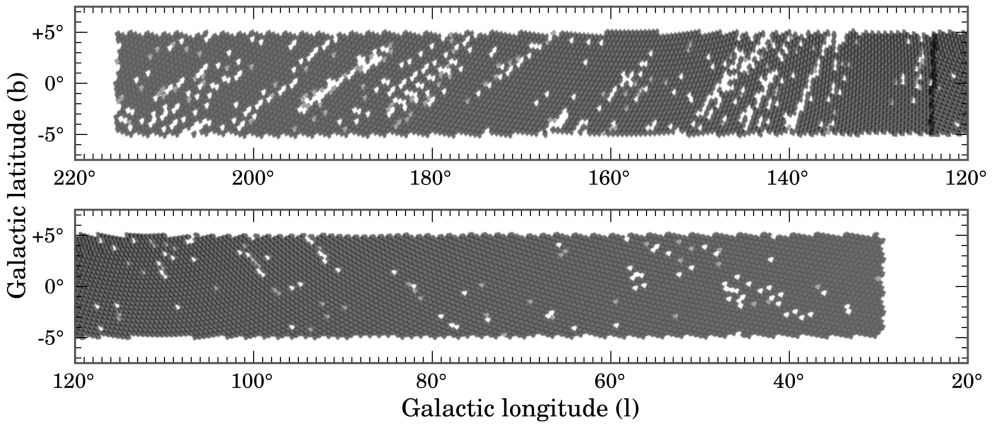

Figure 11: The coverage of the IPHAS survey at the time of DR2. Each semi-transparent point is a quality-approved pointing included in DR2, with small overlaps at the edges of each. Any white gaps are regions where no data has passed the quality control, and the few that are light grey have one offset of the on-off pointing missing. The black strip in the top panel is a side-effect of the tilting pattern which begins at RA = 0h. This figure comes from Barentsen et al. (2014)

.

quality-controlled sets of exposures coming very close to a complete coverage of the northern galactic plane at $> 99\%$ completeness. The majority of the 15270 sets of exposures, around 80% , can benefit from some level of sky correction.

1.6 Data quality

Data released in DR2 (figure 11) was subjected to rigorous photometric tests which set limits on various quality criteria which, if they weren't met, led to data being excluded from the data release and marked as 'bad' data. These criteria are as follows:

- Any data where the 5 σ limiting magnitude is brighter than 20th magnitude in the r-band and 19th magnitude in the i-band and Hα is deemed bad data. The 5 σ limiting magnitude is defined in Barentsen et al. (2014) as the magnitude a point source would have if its flux were equal to five times the level of noise in the sky background.
- The measure of the elongation of the point spread function (PSF) of a point source is its ellipticity and is zero in a perfect noise free environment (where zero means a circular PSF). In DR2, any data exceeding an ellipticity of e=0.3 is excluded.

- Data is accepted in DR2 up to a seeing of 2.5" to maximise the sky area included in the data release. Any images with a seeing greater than 2.5" were excluded from DR2.
- The field pairs of each set of exposures are analysed to check for consistency across the overlapping regions. Any field pairs where more than 2% of the stars had inconsistent measurements of their magnitudes at or above 0.2 mag or where 25% were inconsistent at 0.1 mag were excluded from DR2.
- Every image was visually inspected by a team of 20 members of the IPHAS collaboration looking for extreme effects of clouds and scattered light.
- Source density maps were created which showed the number density of the detectable sources down to 20th mag to check the health of the data across all three bands.
- Only images that contained data across all three filters always taken within 5 minutes of one another were included.

Tests like these are vital in controlling the data since telescope time for IPHAS was allocated for specific dates and not for specific conditions, so a control has to be put in place to limit the flow of poor data into the public domain. In many cases, gaps in the survey were preferred over data of a very poor quality. In the case of this project however, much more of the data is of use since the criteria for imaging data are different to that for photometric data. On average, there are 1.5 observations for each of the 15270 fields, so for the purposes required here, it is almost always better to have some data, than no data and gaps in the mosaics. For 15262 of the fields the 'best' data has been selected and is used in the mosaicking process. Here 'best' data is, in a case where there is more than 1 observation for a field, the data that has the lowest background count levels, which originate primarily from the moon and cloud cover. Only 8 fields of the survey have data where the interference of clouds is so high that they would have a serious detrimental effect on the local region of a mosaic should they be included, and where a replacement observation has not yet been made. For any other fields where only 1 observation has been taken, either because it is of a high enough quality or because they have not yet been repeated, these observations constitute the 'best' data available to us.

Although some of the data not included in DR2 is used for the purposes of this project, we still define 'good' and 'bad' data photometrically, in the same way as Barentsen et al. (2014).

2 The Problem

The background of images taken with ground-based telescopes would, ideally, have zero signal. However, in reality the background of each image varies due to different sources that each contribute varying levels of unwanted light. Understanding each source and its contribution to an image is key to ensuring that accurate readings are taken for any scientific work, as well as for better understanding of the data themselves.

In the IPHAS dataset, a number of fields contain a significant and varying amount of background light, originating from the moon and other dark-sky sources. Here we look at the contributions to the background of the IPHAS data, and the steps taken to understand and account for each.

2.1 'Dark-Sky' sources of background

Airglow is the brightest contributor of light to the dark sky, contributing the main spectral features in the spectrum (figure 12). It originates from emission by atoms and molecules in the upper atmosphere that have been excited by solar UV radiation.

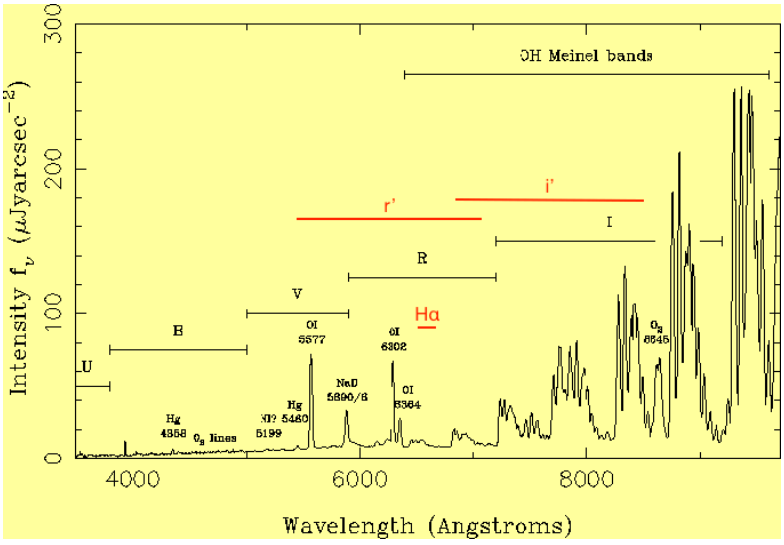

Figure 12: The typical spectrum of the La Palma moonless night sky. The three filters used in IPHAS are marked. The main contributors to airglow in the data are the OI and Na doublet lines (NaD) in r, the O_2 lines in i, and the OH lines across all three. There is a very small contribution from Mercury (Hg) at the shorter end of the r-band. There are also much smaller contributions from other sources, primarily street-light which mostly appears as NaD line emission. Image taken from: www.ing.iac.es/Astronomy/observing/conditions/skybr/skybr.html

The key contributors to the airglow seen in the IPHAS images are lines from hydroxyl (OH, Meinel, 1950b,a), OI seen at 5577, 6300 and 6364 Angstroms (Å), and the sodium doublet (NaD) seen at 5890 and 5895Å coming from nearby street-light on La Palma (Benn and Ellison, 2007). These sources of airglow sit in defined layers in the atmosphere with OH, NaD and the 5577Å OI line at ∼90km, and with the OI 6300Å and 6364Å lines originating from ∼300km.

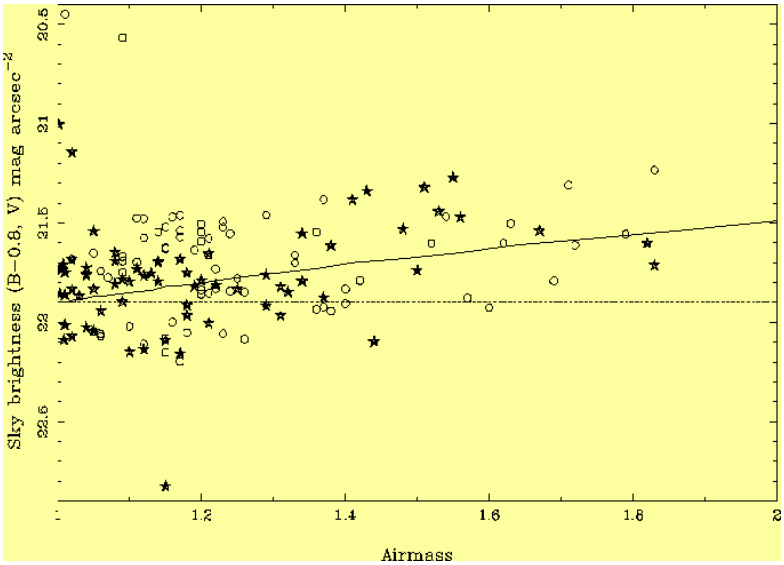

Figure 13: The dependence of sky brightness on airmass at La Palma. Stars represent measurements with ecliptic latitude of >30 degrees, while circles represent <30 degrees. The dashed line represents a typical solar-minimum sky brightness. Image taken from: www.ing.iac.es/Astronomy/observing/conditions/skybr/skybr.html

Airglow varies by 0.03 +- 0.07 mag from the beginning to the end of a night at La Palma (Benn and Ellison, 2007), which is consistent with weak or no variation. However sky brightness is observed to vary with ecliptic latitude by ∼ 0.4 mag between 0 and 90 degrees. Some variation of sky brightness occurs with changing airmass. Airmass is a relative indicator of how much atmosphere the line of sight of an observation passes through. An airmass of 1 means the telescope is pointing straight up. For most IPHAS observations the airmass is between 1-1.5, though very occasionally data were obtained up to an airmass of 2. Unsurprisingly, the sky brightness increases with increasing airmass (so with an increasing atmospheric contribution) as in figure 13.

There is a small (∼<0.1 mag) contribution to the night sky brightness

from light pollution from residents of La Palma itself, as well as the surrounding islands. The majority of contaminating light from human sources is the contribution from the ∼14000 street lamps on La Palma, with a total contribution of ∼0.03 mag in all bands to the zenith continuum of the night sky brightness.

2.2 i-band Fringing

Fringing in the i-band originates from the airglow OH lines seen in Section 2.1 interfering as a result of internal reflections in the CCD chips. It appears as a wave-like structure in the i-band data (figure 14) and contributes ∼2% (Irwin and Lewis, 2001) of the total counts in the images. It is almost entirely removed in the pipeline by CASU, who use a library of i-band fringe frames from other observing runs using the INT WFC to reduce the fringing by at least a factor of 10, though some will remain since it varies significantly on a nightly basis. When i-band data is mosaicked, overlapping fringing is exaggerated in particularly strong cases, and can be visible in a final mosaic as in figure 14. In Fourier space (Section 4), fringing is represented as a quasi-periodic signal, meaning it is not defined as one feature or frequency in Fourier space and as such cannot be removed, without potentially removing astronomical signal as well.

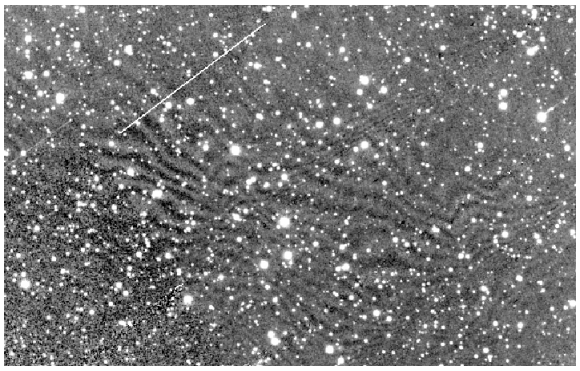

Figure 14: A section of an i-band mosaic with IPHAS data. The wave-like ripples of the fringe pattern are clearly visible through the centre of the image, even after all cleaning and data processing has taken place. There is a satellite trail in this image in the top left.

2.3 Moonlight

The contribution that moonlight makes to IPHAS images is both complex and time variable, and is the main focus of this project. At the beginning of

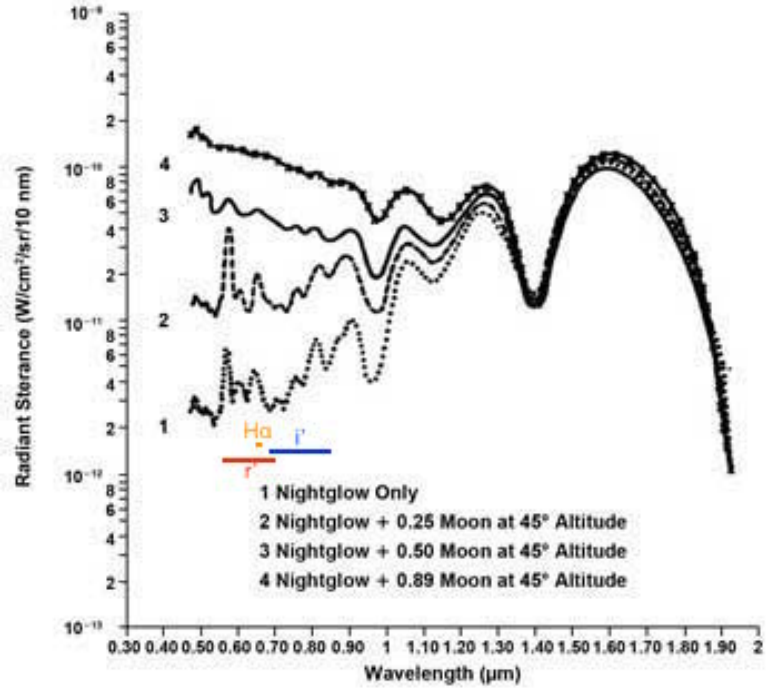

Figure 15: The variation of night glow and moonlight across the spectrum as a function of lunar phase, with the three IPHAS filters marked. The dependence of the strength of the moonlight on lunar phase at the typical wavelengths observed by IPHAS is clear. Image taken from: http : //www.photonics.com/Article.aspx?AID = 50540

data collection in IPHAS, observations were completed at any level of moon brightness throughout the year. Observations during brighter time were, as the analysis went on, seen to have varying levels of background counts in a slope across each CCD. As this factor in the quality of the data became apparent, data collection obeyed tighter rules regarding moon altitude and phase, as well as an increased moon separation (the distance on the sky between the moon's position and the direction the telescope is facing). However the early bright-time data is still the best available for ∼8% of the IPHAS region (at the time of DR2). It is these data, that account for the patches in figure 6. This scattered moonlight must be cleaned up, to allow for future calibration and scientific exploitation of all the data.

The moonlight affecting the IPHAS data is complex in origin. Part of it is straightforwardly-imaged scattered light pervading the night sky. Part of it can also be due to reflection off the inside of the telescope dome and across the CCD array. The result of this is an illumination that is not necessarily uniform across all four CCDs in one observation, and so a fix that works

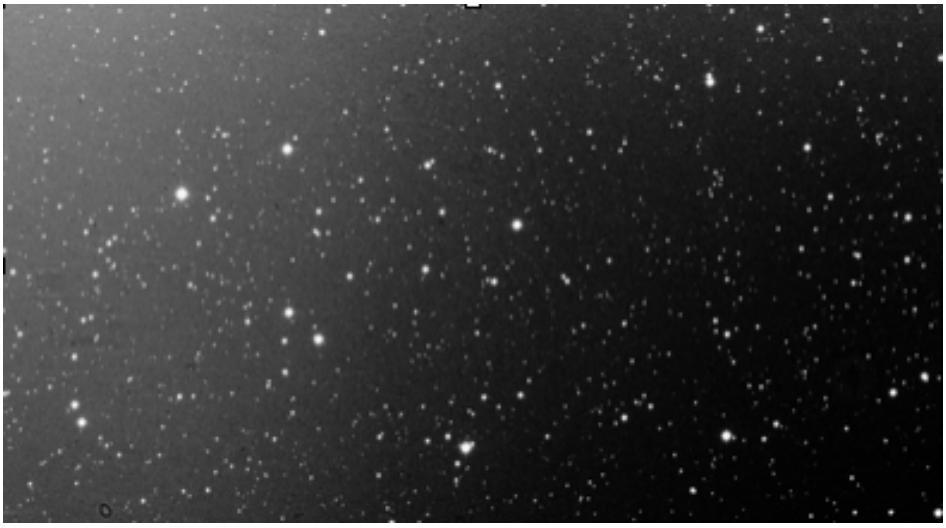

Figure 16: A particularly bad example of scattered moonlight in the background of an IPHAS image.

on a CCD-by-CCD basis is required. It is also noted that cloud cover has an effect on the moonlight present in an image: in a case where the moon is bright yet highly separated from the observed region, any clouds present at the time of observation will alter the scattering of the moonlight across the sky, and may produce a localised gradient of scattered moonlight on the image.

The brightness of the scattered moonlight on the images accordingly varies depending on many factors: the phase of the moon (figure 15), its altitude above the horizon, its separation from the direction the telescope is pointing, as well as the extent (and position) of cloud cover across the sky. Figure 16 shows a particularly strong case of scattered moonlight, which adds anywhere from only a couple of counts to an image, up to the order of $\sim$2000 counts in the worst-case scenarios (these are very rare, and in most cases have been replaced in the winter of 2015).

In order for large scale, fully calibrated and high resolution mosaics to be produced and fully utilised by a survey such as IPHAS, all contributions to the background must be fully understood and accounted for. In the case of IPHAS, the biggest contributor delaying a complete and fully isotropic and homogeneous data set is the moonlight, and understanding and accounting for it in the data set is the focus of this project.

2.4 i or r to clean Hα

The Hα data cannot be directly cleaned of scattered moonlight since this narrowband also contains the diffuse nebulous structure that is of key interest, and it is impossible to separate out the two. Instead, we can make use of the two broad band filters, r and i, and the knowledge of their relative widths to fit a model in one of the broad bands before scaling it to the Hα , thus avoiding removing the counts of interest.

The i-band data has the major advantage of containing no nebulous emission, making it ideal for finding and removing the moonlight background. It does however, as in figures 12 and 15, contain a much brighter dark-time component, predominantly from the OH airglow, leading to fringing.

The r-band data however, while being 'contaminated' by the Hα nebulous emission, has a much weaker and simpler dark-time component that more closely resembles the dark-time component in the Hα narrowband. The r-band data is obtained directly after the Hα , minimising any variation.

The fringing component of the i-band is a complex issue in the cleaning process, but the issue of the nebulous structure in the r-band can be overcome using its matching Hα data, as explored further in Section 5. This, along with the simpler dark-sky component make using the r-band the more ideal path to producing clean Hα data.

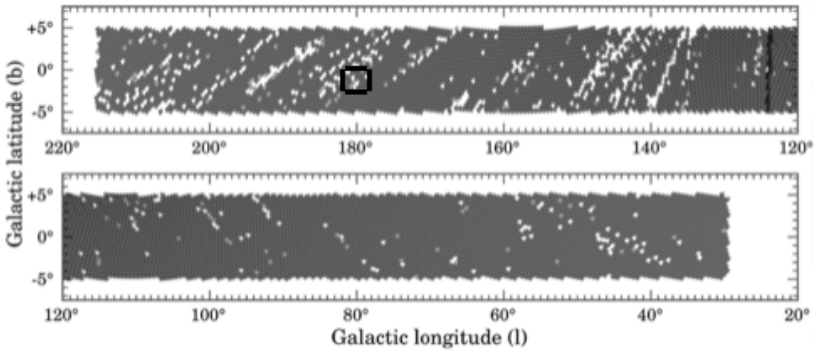

Figure 17: The same figure as figure 11 but with the location of S147 marked as the black square. S147 falls very near the galactic equator at b=-1.6, so has a large amount of surrounding data, as well as falling in a region with a relatively large number of problematic fields- the focus of this project.

Throughout the cleaning process, we use the region around S147 as the test of our method. As well as the details listed in Section 1.3.4, the size and location (figure 17) of S147 make it ideal. Covering $\sim 3^\circ \times 3^\circ$ on the sky at the galactic anti-centre, the images in this region contain a good mix of filamentary structure, high quality data, troublesome fields with a mix of

the problems discussed above, and cover a large region ideal for testing the ability of the method to produce clean, large-scale mosaics.

3 Model Fitting the Background

In order to effectively remove the moonlight (Section 2) and arrive at a flat and dark background, we must, after removing the airglow as accurately as can be expected, have a method to remove the remaining counts in an accurate and adaptable way. A couple of methods are explored in this work. The first, fitting a model to the background for removal, is explored in this chapter.

3.1 Initial data cleaning

Before any model fitting or mosaicking can be attempted on the data, it must first go through some initial data processing. Bad pixels on the CCDs can contribute unwanted counts to the data which must be accounted for, as well as fringing seen in the i-band (Section 2.2) and counts that are associated with the airglow (Section 2.1).

The steps taken are outlined, to account for these issues.

3.1.1 Airglow

Airglow is a significant part of the contribution to the background that must be accounted for before any other processing takes place. To remove airglow from our images, we create dark-time frames as follows:

- A table of dark time frames is formed from the IPHAS catalogue. The images are chosen from good weather nights (seeing < 1.5", no cloud, humidity $< 60\%$ and stable) and occur when the moon is below the horizon.
- The data within this table that is going in to the dark-sky templates is then binned in the same way as the model fitting code (Section 3.2); taking the median of each 100x100 pixel square.
- The median values are collected for each CCD in each filter and the airglow frames are made up of these.
- A model (Section 3.2) is fit to these images in order to produce a smooth template.

These dark-sky templates (figure 18), one for each CCD in each filter (12 in total), are subtracted from the raw data before any other cleaning takes place. Due to the manner by which they are created, they do not contain just airglow, but include other dark-time features, such as fringing in the i-band and background light coming from street lights and nearby settlements.

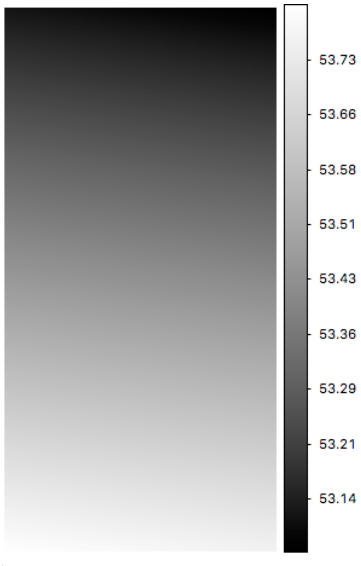

Figure 18: An example dark-sky template. Note that even though it appears there is a gradient across the image, the counts from one end to the other vary by less than 1. This example applies to CCD 1 in the i-band.

3.1.2 Confidence-map cleaning

We must account for standard issues in imaging data that include any known bad pixels on the CCDs, and flat fields which correct the images for errors in the CCD response when there should be no signal. 'Confidence maps' are applied to the data early on in the cleaning process to account for any such camera-based problems.

Twilight flat field exposures are collected every night during observing runs and stacked to form master calibration flats. Providing no instrumental changes, creating the master flats in this way has been found to be stable (Drew et al., 2005). Flat fields contain two components; the gain of an image and its dark-sky response. The dark-sky response is stored as an observation for which there is no input signal entering the telescope. In a perfect world, this would result in a flat and zero response in the resulting image. In the actual case though, the flat field stores whatever the response of the CCDs is when there are no counts to be collected. The gain of the data is the ratio of how the signal given by the CCD varies by how much light there is. The gain should be linear, so as the amount of light coming in to the CCD varies, the response given varies equally. In the IPHAS data collection, the gain differences between each CCD in each filter are found and removed by normalising the average sky in each to CCD 1, something completed by CASU in the pipeline processing.

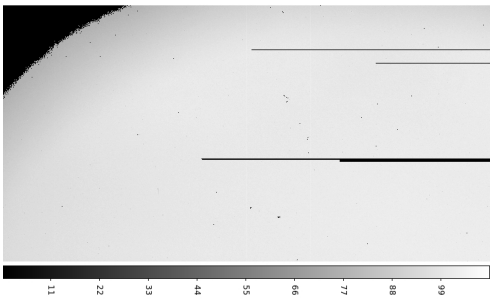

Figure 19: An example confidence map applied to CCD 3 data. The light pixels are high confidence while the dark pixels are low, including the large corner area.

Figure 20: An image from CCD 3 after the confidence map shown in figure 19 have been applied. Note the large area missing in the corner, due to the bad corner on CCD 3.

By combining the master flats; fringe frames for the i-band; and previously known bad pixel maps, an inverse variance weight is calculated for each pixel. The weights make up the confidence maps for each image. These confidence maps give each pixel a 'confidence level', in which the worst pixels have values $\sim$ zero, and the pixels with the highest confidence have values $\sim$ 100. For this work, unless otherwise stated, we have used a confidence level of 90. This means that in a CCD in any given observation, the pixels marked in its confidence map as having a confidence less than 90 are revalued as NaN, while any pixels with a confidence above 90 remain as they are. We have chosen a confidence threshold of 90 in order to remove as much of the unreliable data as possible, without leaving gaps in our mosaics when it comes to stitching all the data together. If the confidence threshold is set too high (e.g. $\sim$ 95) then too much data is masked (particularly in the bad corner of CCD 3) for overlapping data to cover up the gaps, and significant regions of NaN values in the mosaicked images appear.

Figure 21: The custom filter used to mask any last bad pixels in CCD 4 once the confidence cleaning process is complete. This filter has the same dimensions as the CCD images, and the white region is the line of pixels being masked. Black in this image identifies good-quality pixels.

An example confidence map for CCD 3 is displayed in figure 19, with an example image which has been confidence cleaned in figure 20. This figure looks at CCD 3 in particular because it is known to contain the largest region of low confidence data, in its top left corner. The issue of the corner in this CCD is consistent throughout the survey data, however there is enough overlapping data for it not to be a problem.

There is some bad data that the standard confidence maps don't remove, so we constructed our own custom masks to remove, in the same way, any last artefacts found throughout the data. With an example shown in figure 21, these custom masks allow us to tweak the effect of the confidence cleaning step without having to interfere with the confidence masks themselves.

3.1.3 Bright star correction

In some cases when the model is being fit to an image, it is influenced by bright stars and their surrounding scattered light in the image. Especially in cases where the background gradient being removed is subtle (in some cases of the order of a couple of counts) bright stars have a profound influence on what the model fits and later subtracts from the data.

When a bright star falls in one of the CCD images (figure 22), the star not only contributes a high number of counts to the image, but the surrounding sky is much brighter also, due to atmospheric effects and saturation of the CCD. The saturation limit of IPHAS is around the 11th magnitude, so the effect of saturation will play a role for any stars as bright or brighter than this level. When the model (Section 3.2) is fit to the image, it will try to fit all of these extra counts, influencing both the steepness of the slope being fit

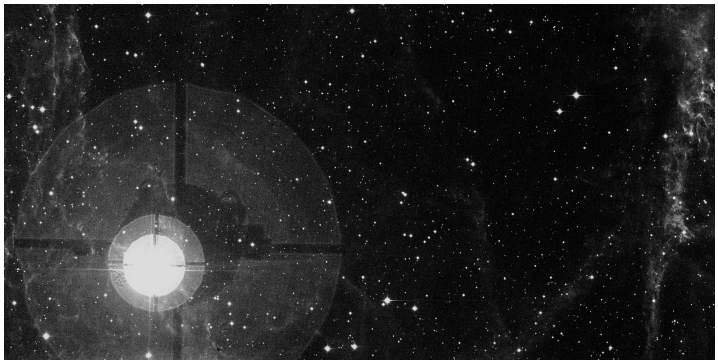

Figure 22: An example image containing a bright star (V-band magnitude of 6.26). The saturation of the star on the CCD is clearly visible, along with its effect of increasing the counts in the nearby pixels, particularly where lens effects come into play. A binned version of this image showing the bright star's effect on the model fit can be found at panel (b) of figure 23.

and its orientation on the image (figure 23).

In order to account for this problem, we begin by finding all cases where a bright star falls on or near an image. For a list of bright stars, we use the Tycho-2 catalogue (Høg et al., 2000) of the 2.5-million brightest stars in our galaxy, 71,307 of which fall within the IPHAS region. The catalogue is 99% complete at V-band magnitude ranges of 0-11. The full magnitude range is used in the cleaning process: stars near the 11th magnitude end of the list have an influence on the model fit when a few are grouped within close proximity of each other on an image.

Listed here are the steps taken to account for bright stars during the cleaning process. The code can flag up any cases where stars fall on or near enough to images that they can have a negative effect on the model fit, and remove the data around them:

- We specify a radius, R_{cut} that is proportional to the magnitude of the star (larger radius for brighter stars and smaller radius for the dimmer stars), and create a list of images which stores any case where a star falls on an image, or where the image edge falls within a distance, R_{cut} of the star.

- We cut out of the images all counts in the pixels within a circle of radius R_{cut} from bright stars, and replace them with NaN values.

- Next, we fit the model via the model-fitting code as in Section 3.2 on the data which now has its brightest stars masked out and, lastly, we subtract the model from the data before bright star removal.

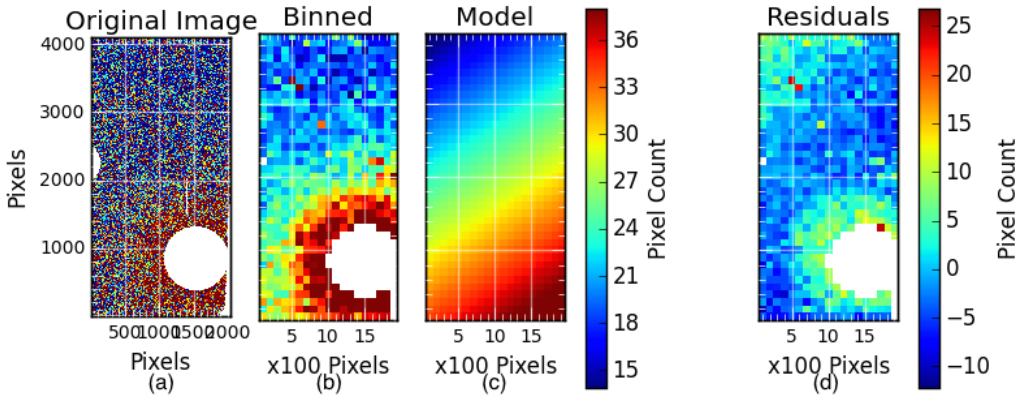

Figure 23: The same star as in figure 22. Here, inadequate masking of the bright star has taken place (panel 1), and the effect it has on the surrounding pixels is clearly visible in the median-binned image in panel 2. In panel 3, the model fit (see Section 3.2) to the data has been affected by the star, resulting in a new gradient appearing in the residuals in panel 4 when the model is subtracted from the image.

The value of R_{cut} is calculated using an equation derived through trial on a broad range of magnitudes of stars across the IPHAS catalogue and is as follows:

$$R_{cut} = \frac{8192}{mag^2} + \frac{1000}{mag} + 100 \tag{1}$$

where mag is the magnitude of the star being masked.

This method of masking means we are able to model fit and subtract the moon light background, without the accuracy of our method being influenced by any bright stars.

3.2 Background model fitting

With our images at a point where all confidence cleaning has taken place, the next step is to remove the remaining gradient left by scattered moonlight. We apply our model fitting process throughout this section and Section 4 on the i-band.

The first step is to bin the images in a way that removes the remaining stars in the images and allows us to measure the underlying background levels. If the bin sizes are too large, we will not be sampling the background in enough detail to pick up the subtlest moonlight gradients, and if they're too small, the remaining stars in the images will impact the values of the bins.

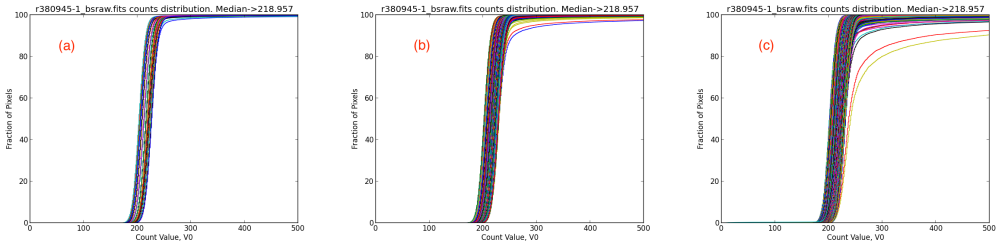

Figure 24: Panel (a), the cumulative distribution of count values (V_0) in each bin in an example image, using a bin size of 400×400 pixels. Panel (b) contains the same for a bin size of 200×200 pixels, while panel (c) uses a bin size of 100×100 pixels. The median count for this image is 219.

The images are split into bins 100 x 100 pixels in size, in a grid pattern with strip of $\sim$ 100 pixels width around the edge to ensure low-quality edge pixels are not included. In each of these 100 x 100 pixel bins, the median value is taken, which leaves a new image containing only background, and no stars (figure 23, panel 2). A bin size of 100x100 pixels is chosen as a compromise between sampling the background in as much detail as possible while not allowing stars to influence the medians in each bin. In the example image used to create these cumulative distributions in figure 24, the typical count levels of the stars are over $\sim$300. Panel (a) uses a bin size of 400x400 (giving a total of 50 bins), and with a median (in this case) of $\sim$219 counts, the median value is not influenced by the stars, which fall in only $\sim$3% of the pixels in the image. In panel (b), with a bin size of 200x200 (giving a total of 200 bins), stars are beginning to have an influence on some bins, though the median, at the 50th percentile, is still a good representation of the non-stellar counts. By panel (c) and a bin size of 100x100 pixels (giving 800 bins), a small number of bins are heavily influenced by the brightest stars in the image and the median in these cases becomes less reliable as representing the background. However, this issue is solved when the brightest stars in the data are masked before the image is binned for model fitting, described in Section 3.1.3. Using a bin size smaller than 100x100 though risks much dimmer stars that are not later masked having an influence on the median as a strong indicator of background counts. The number of stars in an image is not seen to influence the measured sky value in this way.

3.2.1 Fitting the model

Once the images are binned, a representation of only the background remains, leaving us in a position to fit a model to the data. The model we fit is a

linear, two-dimensional gradient of:

$$Z = Ax + By + C \tag{2}$$

where A, B, and C are the free parameters with A describing how the gradient varies in the x direction; B describing how the gradient varies in the y direction; and C being a constant. Z is the counts value in each bin.

Our initial method of fitting equation 2 to the images followed a brute-force Monte Carlo approach, sampling every position in a pre-defined parameter grid for A, B and C until the best fit to the observed values (in this case, Z) was found. Since every value within pre-defined ranges is analysed - for example, between -1 to 1 in steps of 0.1 for A and B, and between 0 to 1000 in steps of 1 for C - the process involves analysing 400,000 different combinations. The starting range of -1 to 1 was selected for A and B because the the variation of the background across either axis never exceeds more that approximately 5% of the constant background level (C). For example if the variation were 5% of C in the x direction, then A would have a value of 0.05. Therefore, setting a starting range of -1 to 1 for A and B means that the best fit will always lie somewhere in the pre-defined parameter space. Next, with a best fit for each parameter found at this level, a new range of values are defined for each at a higher level of detail, usually to another decimal place. In this example case, smaller ranges around the best fit values for A, B and C would be set, and would be analysed in steps of 0.01 for A and B, and 0.1 for C until best fit values for all three are found at a satisfactory sensitivity.

A major problem with this method was the time it took to run, and the tendency of the fitting method to fixate on local minima. Even increasing the detail in one variable, for example steps of 0.01 in A, increased the number of combinations being analysed by a factor of 10. The second problem, the finding of local minima becomes a problem since the method began by looking at the parameter space broken up into course steps, before focussing in on good fit values in more detail. In the case of our model, the values found as best fitting A, B and C at this point may not necessarily be the best fit to the data, but instead may be a local minimum in the parameter space, where the model fits somewhat, but is not the global minimum that is the best fit to the background gradient.

3.2.2 Markov Chain Monte Carlo fitting process

To combat these problems of local minima and slow-processing times, we instead use a Markov Chain Monte Carlo (MCMC) fitting method. MCMC model fitting still fits the same linear gradient model in equation 2 but does

not explore the parameter space in a fixed way, rather in a randomised way that prioritises areas of the parameter space that provide a good fit.

An MCMC code begins by picking a random point within the parameter space and calculating the quality of the fit at that point. A second random point is then chosen, and the fit at this second point is compared to the first to see if it is better or worse. The third point is then chosen, which is weighted towards the best fit by what the code has learned from the first two points. This continues for a pre-defined number of 'burn runs' - runs that help the MCMC code locate the general area of the best fit. The movement of the code from one point to the next is called a 'walker'.

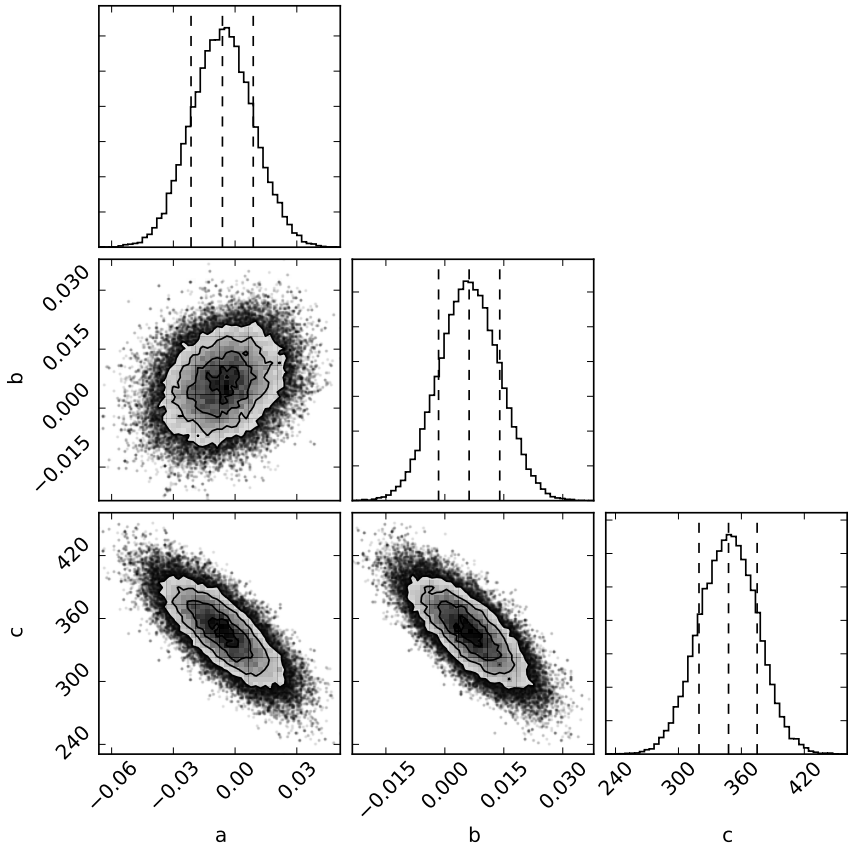

Figure 25: The triangle plot produced by the MCMC model fitting showing the explored parameter space for each variable (A, B and C). The three histograms show the distribution of likelihoods resulting from the MCMC code, along with the median value as well as the 16th and 84th percentiles marked as the dashed lines. The three density plots show the likelihood distribution in two dimensions for each available combination.

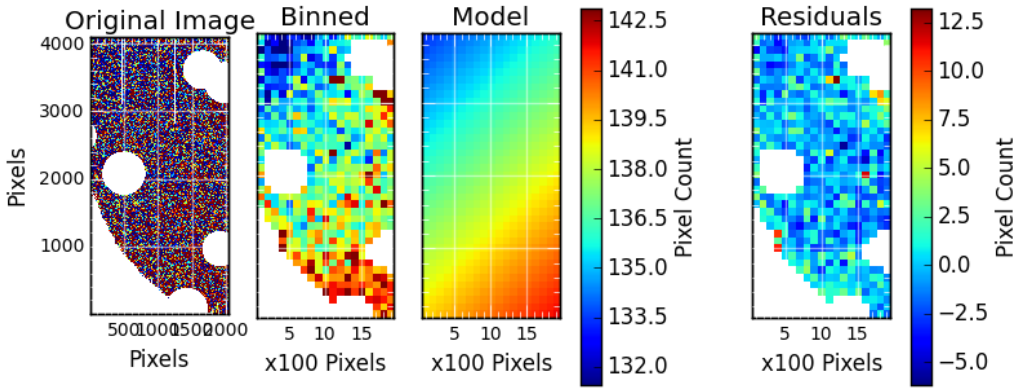

Figure 26: An example image that has passed through the MCMC model fitting process. In panels 1, 2 and 4 the areas where bright stars have been masked (Section 3.1.3) appear as white circles. Panel 2 shows a gradient when the image in panel 1 is binned, which is modelled well in panel 3. The residuals after model subtraction in panel 4 show a much flatter background, with counts $\simeq$ 0.

After the 'burn runs' are complete, the walker continues exploring the parameter space around the now-localised best fit for another set of predetermined runs. For all of these runs, the values of A, B and C trialled by the walker are stored in a table (the 'burn runs' having been discarded or 'burnt'). Since the MCMC code has already honed in on the general area of the best fit in the parameter space, the points that the code trials a fit on once the burn runs are complete are the same as the distribution of likelihoods of the best fit. These distributions of likelihood are plotted up as 'triangle plots', an example is seen in figure 25.

We use an 'ensemble' sampler code written in Python, called EMCEE (Foreman-Mackey et al., 2013). An 'ensemble' sampler code uses a specified number of 'walkers' to explore the parameter space multiple times at once, with a random starting value for each. The walkers communicate this information with each other at each step, making this method very efficient.

In the case of this work, I have used 100 walkers completing 200 burn runs, followed by a further 1000 runs. This number of burn runs is enough that the parameter space has been sufficiently sampled and the general region of the best fit found. Using 100 walkers for 1000 further runs allows the code to be quick yet accurate by essentially computing 100,000 runs, but by 100 walkers simultaneously.

The distribution of likelihoods of the best fit are shown in figure 25. The median value for each parameter from these distributions is taken as the best fit.

Uncertainties on A, B and C are taken from the 16th and 84th percentiles of the distribution of values in each parameter that the MCMC code has explored. Our uncertainties then become the $84th - 50th$ percentiles for the upper bound and the $50th - 16th$ percentiles for the lower (since the 50th percentile is the median which we take as our final value). Since, as we see in figure 25, the distribution of the model fits inside the parameter space appears symmetric, we simplify our uncertainties to one value, of:

$$\sigma = (84th - 16th)/2 \tag{3}$$

Figure 27: The model that is fit in figure 26. Once A, B and C are known, the model can be reproduced in the same dimensions as the original image, before being subtracted off.

Once the best fit values of A, B and C for each image are found, we are able to recreate the model at the scale of the original image using equation 2. As seen in figure 26, this model (figure 27) is directly subtracted away from the image (after confidence cleaning, but before any bright star masking) leaving a flat residual background with a median of ~ 0.

To judge the quality of the model fit on the data, the root mean square (RMS) residuals after the model has been subtracted from the image are calculated and compared with the residuals if a flat (no gradient) model was used. The RMS tells us how much of the large-scale background gradient remains unfit and so gives an indication of how well the model has fit the moonlight. The RMS residual for an image is given by:

$$RMS = \sqrt{\frac{\sum_{i=1}^{n}(image - model)_i^2}{n}} \tag{4}$$

where n is the number of bins in the image and $image - model$ is the value of each bin in the image with the model subtracted.

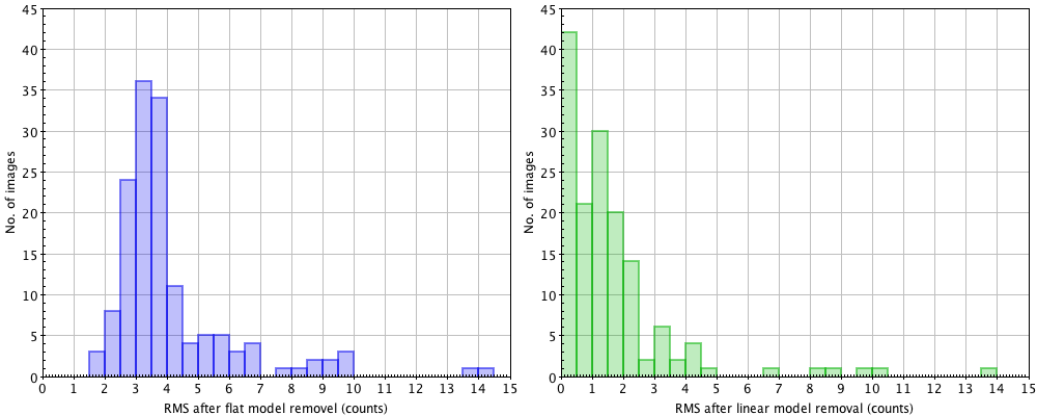

Figure 28: (Left) the RMS values of each image in a patch of sky in S147 (containing 150 CCD images) with a flat model removal and (right) the same images with the linear model removed.

In figure 28 there is a shift of the RMS values towards zero from the results of a flat model removal in the left-hand panel (where no images have an RMS less than 1.5) to the results of a linear model removal on the same data in the right-hand panel. A gradient model is a much-improved fit where, in most cases, the majority of scattered light is removed, with some outliers remaining.

3.2.3 Alternative models

As well as the linear two-dimensional model, other models were trialled to see whether a more complex fit would have an impact on the quality of the clean on the images. Figure 29 displays the results of such testing. The first alternative model trialled was a radial gradient given by:

$$A.r + B = Z \tag{5}$$

where A is the variation with increasing radius, r, and B is constant. The radial model includes two other parameters, (C, D), where C and D are the pixel coordinates where the centre of the radial function lies. We see that for a radial image, in figure 29 panel (c), the RMS values of the resultant images have only little variation from those of a linear model (panel a).

The other alternative model applied was a quadratic:

$$A.x^2 + B.x + C.y^2 + D.y + E = Z \tag{6}$$

where $A.x^2 + B.x$ is the quadratic function fit to the x component of the background gradient, $C.y^2 + D.y$ is the quadratic function fit to the y com-

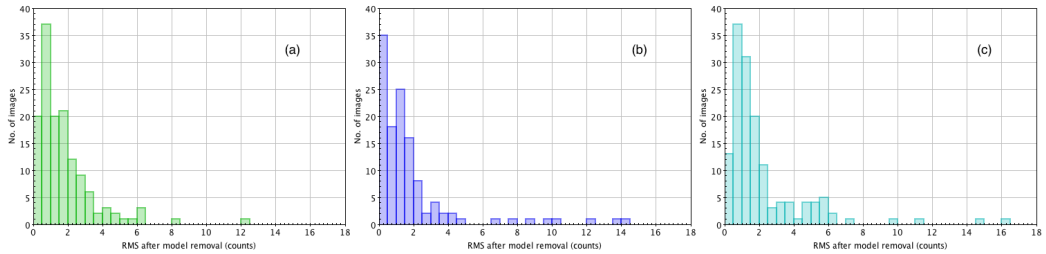

Figure 29: The results of applying different models to the same patch of sky used in figure 28. Panel (a) contains the RMS values for the linear model for 150 images. Panel (b) contains the RMS values for the same 150 images using a quadratic model, and panel (c) using a radial model.

ponent of the background gradient, and E is the constant. This model was applied in the same way as the linear model. Since the MCMC code had to find the best fit to 5 parameters compared to the 3 of the linear model, this model took longer to compute. In figure 29 panel (b) we see that, although there appears to be some improvement with this model, there is not a large enough improvement to say for sure that the more complex model, with its much longer processing times is a better option over the linear model.

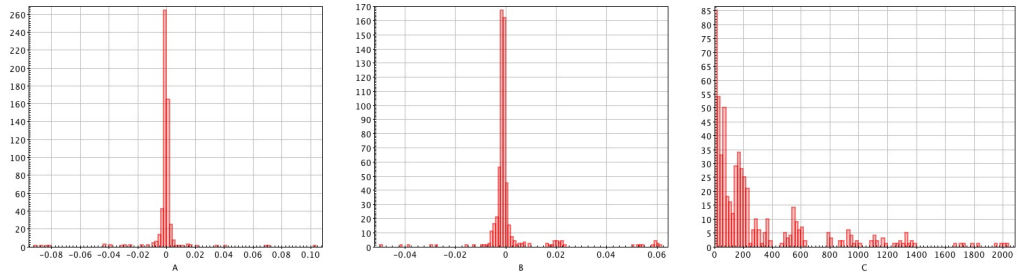

Figure 30: The spread of the A, B and C parameters found for each image in a selected patch of sky. Both A and B peak at around zero, which is expected as a perfectly flat background would have A and B of zero, and the majority of cases of moonlight have subtle gradients. C describes the overall brightness of the moonlight in the image backgrounds, which is generally quite small, but includes some cases where the moonlight is considerably bright (high moon phase, high moon altitude).

A comprehensive and simple linear model has been shown to fit and remove almost all scattered background light in the images, though fringing will still remain. The data has been cleaned of bad pixels and dark-sky counts, before being binned up to remove the stars and create an image representative of only the remaining background levels. The best fit for the model is found for these binned values, before being subtracted straight out of the data to produce a cleaned image. A linear fit is chosen since more

complex models take a longer processing time and become unnecessary as there is no significant change in the resultant images and their RMS values. This model fit (with example results in figure 30) provides the core of the final cleaning method to produce clean Hα mosaics, though in the final process the i-band is not used. Due to a number of factors including the variable fringing in the i-band, and that we have the Hα data to account for nebulosity, the r-band data is used to reach our eventual goal, with the process described in detail in Section 5.

3.2.4 Iterative approach

In the final version of the code, an iterative method was implemented. In this approach, the model fitting is completed as normal, the difference coming when the model is applied. When the fitted model is removed from its corresponding image, the residual 'cleaned' image is then binned to reveal only the residuals in the background, and passed through the model fitting process again.

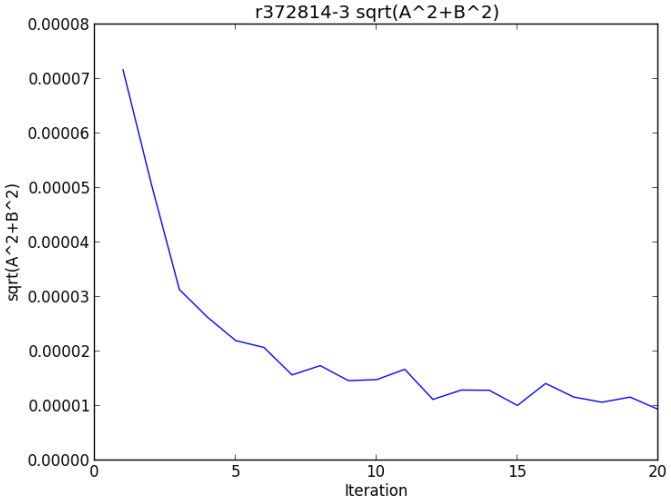

Figure 31: The convergence of A and B in the iterative model fit of an example image. Over iterations, the model converges close to zero.

The convergence of A and B in the model fit over multiple iterations is plotted in figure 31. If A and B converged to zero, there would be no background gradient of any kind in the image (the ideal case) however it is impossible to remove absolutely everything. Running multiple iterations on each image though has a significant increase on the quality of the fit.

In the final model fitting code, we use three iterations of model fitting and removing from each image. Three iterations is chosen as a strong middle ground between a more accurate and comprehensive background cleaning avoiding more complex models, without drastically increasing the computation time of the code to a point where it becomes impractical to run on large areas of sky.

4 Fourier Fitting

In this chapter I present my exploration of Fourier Transform (FT) methods to remove the scattered light gradient from the IPHAS i-band images, including an introduction to the Fourier Transform process and a discussion of its limitations when applied to our data. The principal idea behind trying this method was to be able to isolate a single feature or set of features in Fourier space - by Fourier transforming the IPHAS data - that corresponded to the scattered moonlight. Once identified, it would then be a case of masking those features within Fourier space before translating the data back to the spatial domain. The resultant image would then no longer contain its scattered moonlight component, and would therefore remove the need for heavy processing and could, potentially, be highly accurate.

Joseph Fourier laid the foundations to the concept of splitting a function of a variable into a series of sine waves of multiples of the variable. Developed from this concept, and dedicated to Joseph Fourier, Fourier transforms split a signal into its various frequency components, an idea which can be applied to imaging data. The Fourier transform of an image decomposes the image into the spatial frequencies that it is made up of. Since a background gradient in an image will be associated with a specific set of frequencies, in principle it may be possible to use the Fourier transform of IPHAS image data to identify and remove this gradient in Fourier space.

When an image is Fourier transformed it moves from being represented in the spatial domain, to its equivalent representation in the Fourier domain, otherwise known as the 'frequency domain'. In the frequency domain, each individual pixel represents a frequency contained in the whole spatial domain image.

4.1 Practical implementation of Fourier image analysis

When applying Fourier transforms to images, as in the case of this work, not all of the frequencies that form the image in the spatial domain are contained in the equivalent Fourier domain image. Instead a sample set of frequencies is used, equal to the size of the spatial domain image, that is large enough to completely describe it. This is known as a 'Discrete Fourier Transform' (DFT, Bracewell, 1999) and results in the Fourier domain image being of the same size as the spatial domain image. A Fast Fourier Transform (FFT) (Bracewell, 1999) (Brigham, 1988), as used in this work, is the name of one of a set of algorithms that compute the DFT, but require much less computation to achieve the same result. An FFT method then computes the DFT much faster than computing the DFT directly. A DFT has to compute

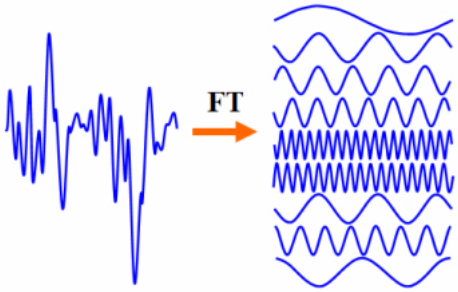

Figure 32: A Discrete Fourier Transform (DFT) takes a complex signal, like an pixel in the spatial domain, and splits that signal into sinusoidal waves of various frequencies, like a spectrum (as shown here) or an image in the Fourier domain. Image taken from: www.revisemri.com/questions/creating_an_image/fourier_transform

N^2 computations, where N is the number of points (or 'signals'). An FFT splits that signal into smaller components, for example into two signals of length $N/2$. Computing the DFT then on one of these two smaller signals will take $(N/2)^2$ computations, meaning $2 \times (N/2)^2$ computations for both signals. This simplifies to $(N^2)/2$ which is half as many computations as in the DFT. A FFT actually splits the signals into even smaller components, such that the number of computations needed to create the Fourier domain image is $N \times logN$.

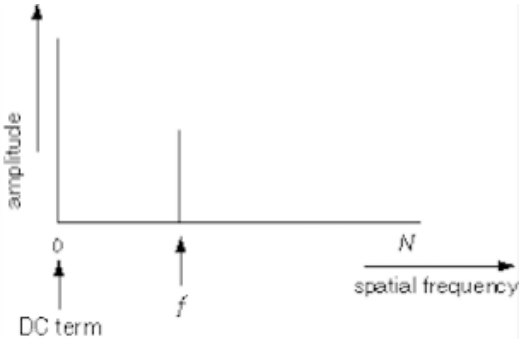

Figure 33: In this plot, N, along the x-axis, is the spatial frequency, and the amplitude of each frequency is plotted in the y-axis. The signal here contains only one spatial frequency, f. The zero-point frequency (or 'DC term') is the sum of all the frequencies making up an image. Since the DC term here is zero, the image has an average brightness of zero, which implies that f alternated between equal positive and negative values in the spatial image. Figure taken from: cns − alumni.bu.edu/ slehar/fourier/fourier.html

The DFT takes an image in the spatial domain and splits it into the

separate sinusoids that describe the signal seen in the image. By the time the image in Fourier space is produced, the two dimensional Fourier transform has completed a one dimensional Fourier transform (figure 32) on every row of the image in the spatial domain, and again on every column, giving a full 2-D Fourier transform with the same size as the original image. The resulting Fourier domain image stores the frequency, amplitude and phase of each sinusoidal wave, ranging from the zero-point frequency (ZPF), which stores the average brightness in the image, to the N-1th frequency, decided by the number of pixels in the original, spatial image. Each pixel in Fourier space represents one spatial frequency, while the signal of each pixel represents the amplitude of the frequency represented at that point. The ZPF is otherwise referred to as the 'DC term', after its equivalent in an electrical context, in which the DC (Direct Current) is the 0Hz term (our ZPF) while the AC (Alternating Current) components are equivalent to the non-zero bits in our Fourier domain image (all frequencies except the ZPF).

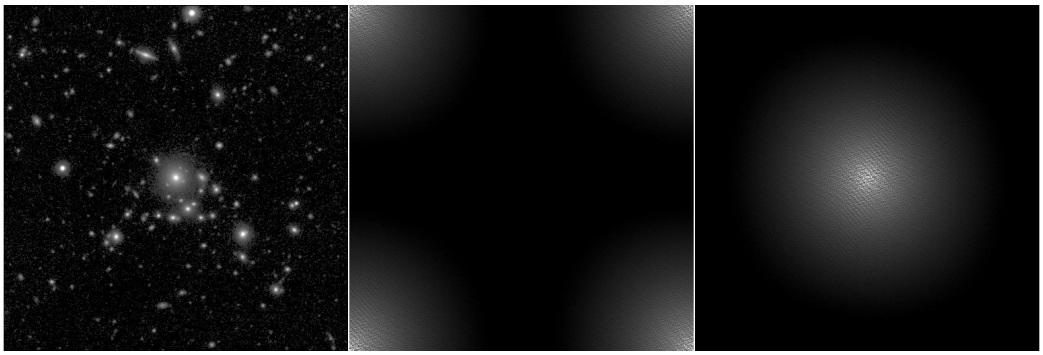

Figure 34: (left) The raw image in the spatial domain with a flat background used in the example case, (centre) the Fourier transform (frequency domain image) of the raw image and (right) the sorted Fourier transform.

In a spatial domain image with a range of positive and negative counts, the ZPF will give the sum of the frequencies in the Fourier transform. Therefore, in an image where the signal (total number of counts) below zero is equal to the signal above zero, the ZPF will have a value of 0 (figure 33). In real images however, it is impossible to have a count that is negative. This means that in the case of the IPHAS data, this ZPF value will be the sum of only positive values, and so will give the sum of all the frequencies in the Fourier transform of an image.

The function for the 2-D DFT for each point in a Fourier domain image ($F(u, v)$) is as follows:

$$F(u, v) = \sum f(x, y) e^{(-2\pi i \frac{(ux+vy)}{Nx,y})} \tag{7}$$

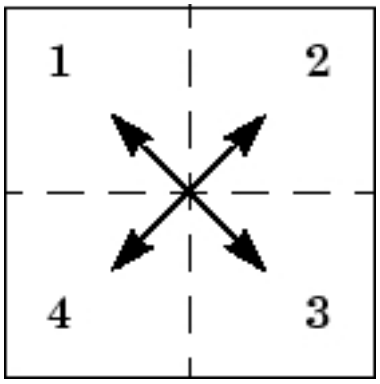

Figure 35: When the Fourier domain image is sorted, quadrant 1 swaps with 3 and quadrant 2 swaps with 4 moving the ZPF to the centre of the image and sorting the high frequencies to the centre and the low frequencies to the edge. Image taken from: uk.mathworks.com/help/matlab/ref/fftshift.html

$$f(x,y) = \sum F(u,v)e^{(+2\pi i \frac{(ux+vy)}{N_{x,y}})} \tag{8}$$

where $f(x,y)$ is the image in the spatial domain, $i = \sqrt{-1}$, and the exponential is the base function that corresponds to each point in Fourier space. $N_{(x,y)}$ is the length of the x and y axes. During the Fourier transform, the sum is first run on each row in the image data, treating it as a 1-D Fourier transform and splitting it into its component sinusoidal waves of varying frequencies, as in figure 32. The process is then repeated for every column in the image with each frequency fitted and stored in the Fourier domain image- with lowest frequency at the centre and highest at the edge. The reverse Fourier transform equation (equation 8) is also given, and simply reverses the process to convert from the Fourier domain back to the spatial domain.

We can see from equation 7 then, at the ZPF:

$$F(0,0) = \sum f(x,y) \tag{9}$$

i.e. the zeroth frequency is the sum of the spatial domain image.

Once an image has been converted from the spatial domain into the Fourier domain (figure 34, panels 1 and 2), it is useful to rearrange the Fourier image to have the ZPF at the centre for visualisation and manipulation. The result of this sorting (figure 35) leaves the high-frequency components near the centre of the Fourier image, and the low frequency components at the edge. Sorting the Fourier domain image leaves it radially symmetric and allows us to both plot its power spectrum as a function of radius from the centre, and create simple filters to mask the data. Small objects and defined edges in an image are described by the high frequencies now found at the

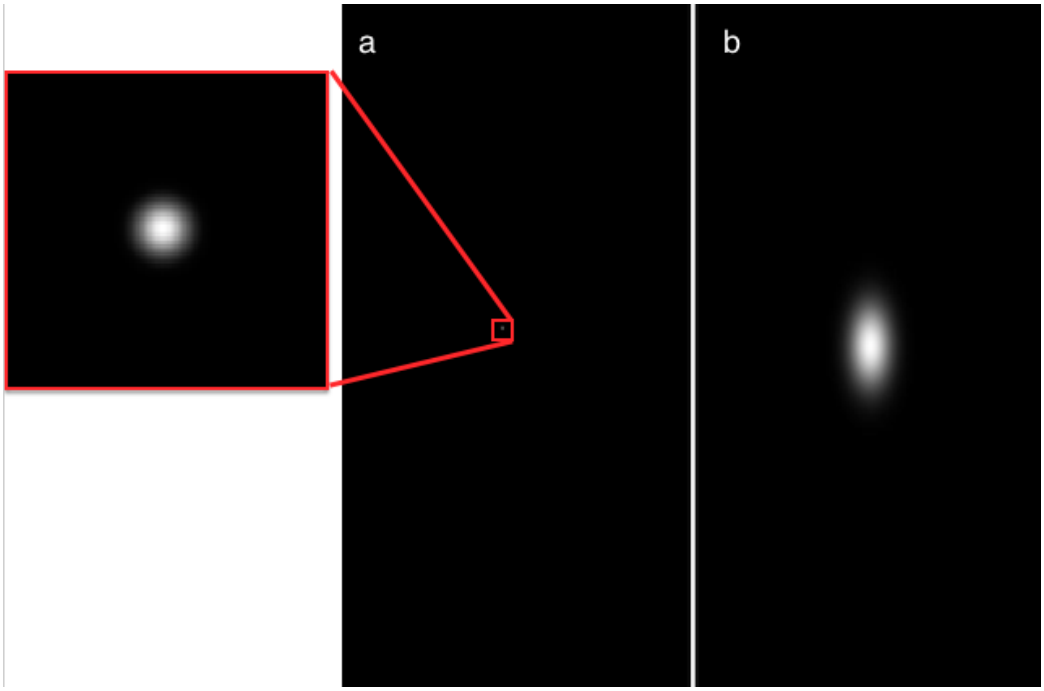

Figure 36: The image here (a) and its Fourier domain counterpart (b) contains a single off-centre star with a Gaussian profile.

centre of the image in Fourier space, while large-scale, diffuse objects (e.g. nebulosity, sky background) are represented by the low frequencies located at the edges.

4.1.1 Features in Fourier space

In order to know which data in Fourier space should be masked in order to remove the scattered light background, we must first explore how features in the spatial domain translate into a Fourier domain image.

The first image, figure 36, is the FFT of an image with a background of 0, containing one star with a Gaussian profile off the centre of the image. The result of this image in Fourier space bears strong similarity to that of many stars in figure 37, except the 'bright' central region contains much smaller values and doesn't extend as far. This suggests that the values in Fourier space stack for each feature in an image, and that stars are described by a whole range of frequencies.

The second image, in figure 37, shows the transform of an image containing only stars and background of zero. The result is an FFT that, when sorted, contains the majority of its information at the centre of the image.

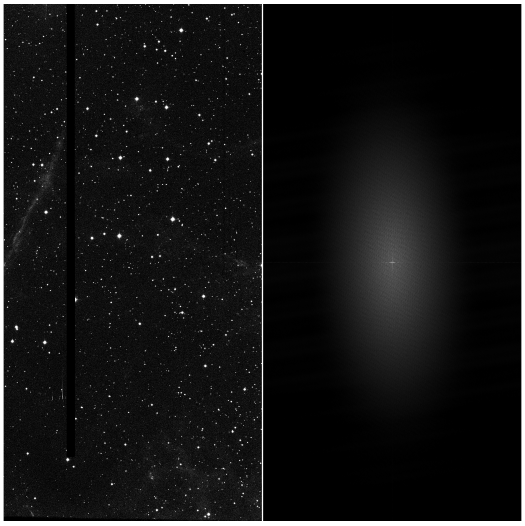

Figure 37: The image displayed here (left, an IPHAS image) and its Fourier transform contain many stars. The Fourier domain image structurally looks similar to the case of one star in figure 36.

The transform shows a profile that is brightest at the centre, and decreases in brightness with increasing radius.

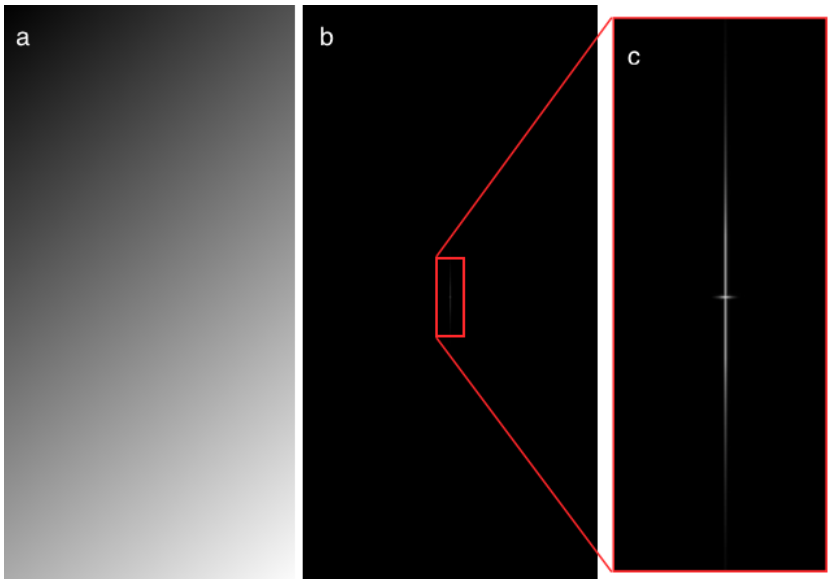

Figure 38: The image here contains only a gradient across its axes (a) which, when Fourier transformed (b), is translated into a one pixel-wide cross (c) running through the centre of the Fourier domain image.

The third example, figure 38, contains an image containing only a background gradient, with no other features. The gradient translates into Fourier space as a line with a width of 1 pixel passing through the centre of the image. In fact the background forms a cross pattern, with the x and y components of the cross each representing the x and y components of the gradient. In this case only that component is represented in Fourier space. The power spectrum of one direction of this cross shows the strong peak at the centre of the transform, but also that the cross extends all the way to the edges of the image, albeit at a relatively small magnitude. Therefore the gradient in the image is represented by a large range of frequencies, peaking highest at the centre with the ZPF, but with some small component visible at the low frequency end. Herein lies the problem (discussed in Section 4.2.1)- gentle, even slopes are 'broad band spectrum' in Fourier space, making them hard to filter using low- or high-pass filter techniques.

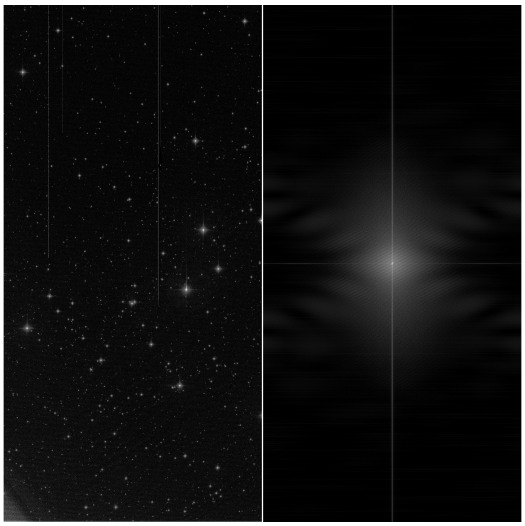

Figure 39: An example of an image from the IPHAS catalogue that is known to have a background gradient and its FFT. This image contains features seen in both figures 37 and 38 with a mixture of stars and background gradient, features which are seen also in the Fourier domain image with the bright centre and cross pattern.

Figure 39 contains the Fourier transform of an IPHAS image known to contain a strong background gradient due to moonlight. As this image is taken in the i-band, it contains only the scattered moonlight gradient, without extended emission in Hα . In its Fourier transform, there are elements that have appeared in the example cases so far, with the bright region corresponding to the stars in the images, and the cross shape through the middle of the right panel of figure 39 representing the unwanted background slope.

This cross contains the information that is responsible for the background gradient we are aiming to remove, so it is this area in Fourier space that must be focused on.

4.1.2 Masking in Fourier space

Here we discuss efforts to mask this cross pattern in the FFT with the end result of having an image free from scattered light, using an example case kindly provided by J. Geach.

- Beginning with a raw image of 512x512 pixels containing galaxies of various sizes and a flat background, seen in figure 34, along with the unsorted and sorted Fourier transforms of the image.
- A linear 1-D ramp is then artificially added to this image along one of the axes, with the resulting image seen in figure 40 panel (a). The image, now with a gradient applied, is then Fourier transformed, seen in figure 40 panel (b). The Fourier transform of the ramp that is applied to the image in panel (a) can be seen in panel (c).
- The gradient that we wish to mask is described by the one pixel wide cross passing through the exact centre of the sorted FFT which peaks in the centre, as seen in Section 4.1.1. Therefore, the mask will only edit data in these few pixels. The mask, seen in figure 40 panel (d), is directly subtracted from the sorted FFT of the raw image.
- The resulting FFT is then inverse Fourier transformed, returning it to a spatial domain image and with the unwanted background gradient now completely removed (figure 40) panel (e).

4.2 Application to IPHAS data

We know, from the previous two sections, that the information relating to the unwanted moonlight gradients lies in the bright cross seen falling on the centre of the reorganised image in Fourier space (figure 27) and peaking with the ZPF at the centre. Since all of the information for the gradient lies here, this is the area to mask. An example where the original model fitting method has produced a good, clean image can be used to get an idea of the shape of the mask required. Shown in figure 41 is the power spectrum of one such case where the model fitting code produced a good fit (figure 27). As is expected, the model of the scattered light background is represented in the cross-like pattern through the centre of the image in Fourier space.

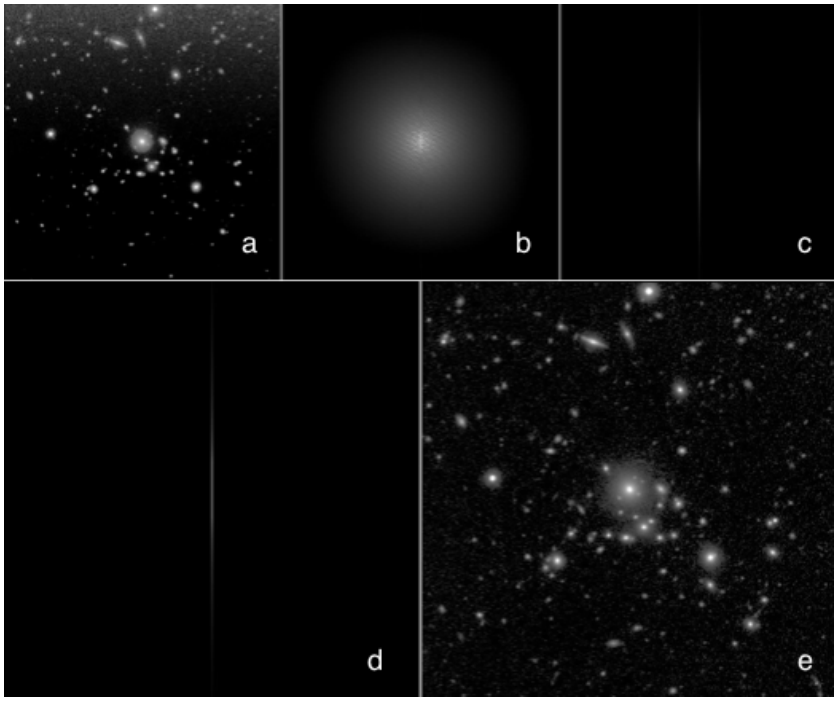

Figure 40: The image (a) with an artificial linear gradient (ramp) applied along the y-axis and its Fourier transform (b), as well as the FT of just the ramp (c). The mask applied to the image is in panel (d) (the same as the FT of the artificial ramp in panel (c), in this case), and the data after having been masked in Fourier space, inverse Fourier transformed to the spatial domain (e). The image in panel (e) contains no residual background gradient (this returns the image back to its original state in figure 34).

In order to remove the scattered light in Fourier space without prior knowledge of its parameters, a Hanning profile is chosen as our filter (Harris, 1978). The Hanning profile seen in figure 42, named after Julius von Hann and otherwise known as the 'Cosine Bell', allows us to account for the scattered light by masking most heavily towards the bright centre of the Fourier transformed image and less heavily towards the edges, according to the shape of the power spectrum of the Fourier transform for a background gradient in figure 41. The Hanning window is defined as:

$$w(n) = 0.5 - 0.5\cos\left(\frac{2\pi n}{M-1}\right) \qquad 0 \leq n \leq M-1 \tag{10}$$

where M is the width of the window function (Harris, F.J., 1978).

Using this profile, a 2-D filter the same size as the IPHAS images is created (figure 43). The mask applied used the inverse of the Hanning profile, as the

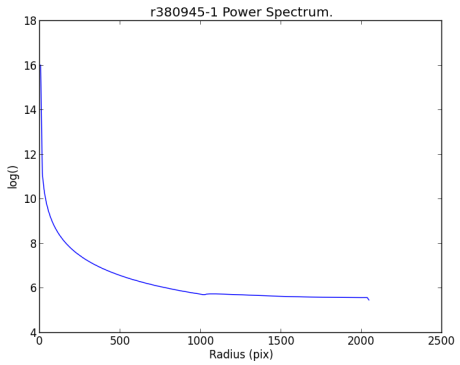

Figure 41: The power spectrum of figure 38 panel (b), displaying the bright peak at the centre.

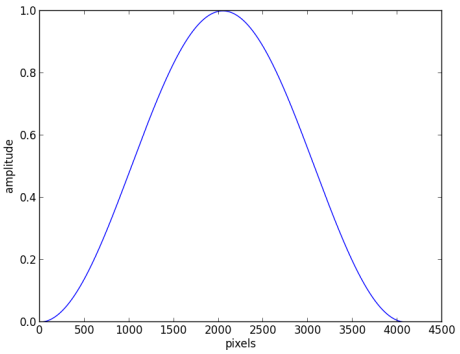

Figure 42: The shape of the Hanning profile of 'cosine bell'. Its key features are its peak at 1, tapering off to 0 at either end. This example extends the y-axis length of an IPHAS image, 4096 pixels.

intended result was to mask most heavily the values towards the centre of the Fourier transformed images. Values inside the mask range from 0-1, with the most heavily masked pixels at the centre having a value nearing 0, and those outside of the Hanning window having a value of 1. To apply the mask, we multiplied it by the Fourier transform of an IPHAS image, then reverse Fourier transform to produce the now masked data in real space, seen in figure 44.

4.2.1 Artefacts

Once the masked data had been inverse Fourier transformed, there existed a large number of bright streaks, or 'artefacts' (figure 44). The cross was only

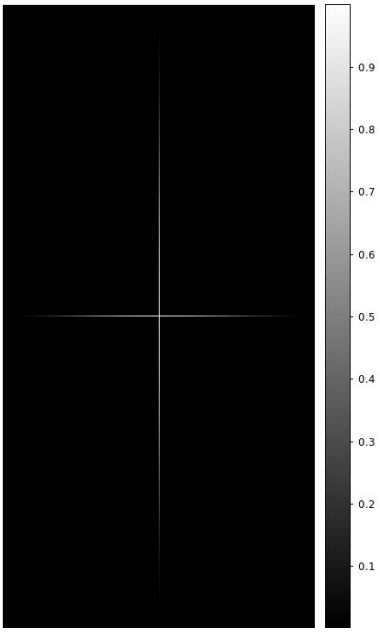

Figure 43: The 2-D Hanning image used to mask the data. Before application, the values in the mask were first inverted, before multiplying the image data by it.

one pixel wide and the strongest masking happened in the central, brightest pixels in the Fourier space image. Since the cross was very narrow the mask did not apply a smooth ramp between masked and unmasked values. The result of this type of 'simple masking', where heavily masked and unmasked pixels neighbour one another, gave a number of artefacts in the newly masked image.

In an attempt to combat this, a new style of mask was created where the counts in the surrounding pixels were reduced by half as much as the central masked pixels. The key difference between this and the previous masking attempt is that streaks across the resulting images are shortened, although not entirely removed.

It is inferred that with a smooth enough taper from the masked pixels to the surrounding pixels these artefacts might disappear. An example of the mask including the buffer zone is displayed in figure 45. This new mask, while still applying a Hanning profile to the cross through the centre of the image in Fourier space, also applies the same Hanning profile in the directions perpendicular to the cross at each masked pixel. So each masked pixel in the central cross has a buffer zone around it in the shape of a Hanning profile whose peak lies on the cross.

The result of these new adaptive buffering masks is displayed in figure 45.

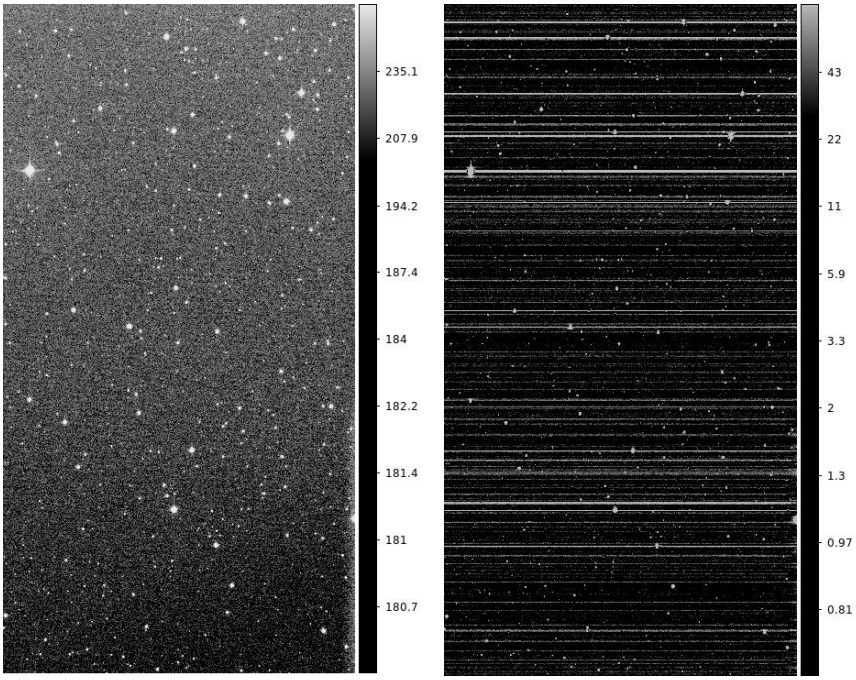

Figure 44: (Right) The resulting image after the mask in figure 43 has been applied in Fourier space and (left) the original image with no masking. After masking, large streaks remain in the resulting spatial domain image.

In this example again, while the artefacts left behind have reduced in magnitude, they are still present. By masking the data to remove the background gradient without knowing the parameters of that background beforehand, we are creating artefacts in the image. In order to remove these artefacts, the pixels around the cross containing the gradient's information in Fourier space must also be masked. As seen in figures 36 and 37, this central region in Fourier space contains the information relating to the stars in the image, and by extending our mask to mask not only those describing the background gradient but these pixels also, the mask has changed the amplitudes of the frequencies describing the stars. These alterations translate into artefacts on the images with, as is visible in the reduced case in figure 45 panel (d), the brightest stars in the image at their centres.

4.3 Critique

We explored converting our images with a background of scattered moonlight into Fourier space as a way of isolating and removing the component associated with the scattered light background. While this method works in theory,

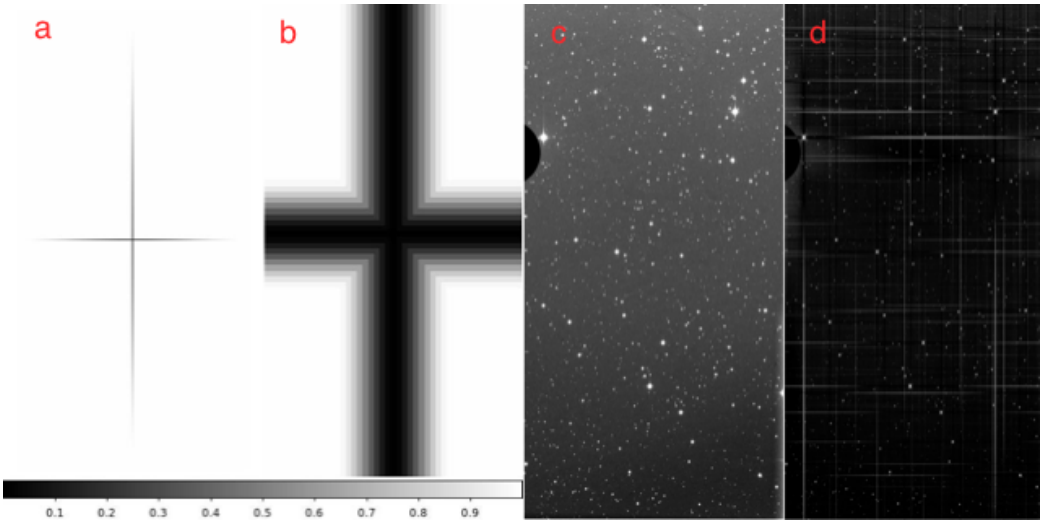

Figure 45: The final iteration of the filter (a) used to mask the cross through the centre of the Fourier domain images to remove the background gradients and a zoom in on the centre of the filter (b). The strongest masking happens at the centre, where pixels are darkest, and reduces with increasing distance from the centre and with increasing distance from the one-pixel-wide cross through the middle of the mask. The original image with background gradient before masking in the Fourier domain is in panel (c) and the same image after applying the mask in figure 45 is in panel (d). The streaks seen in figure 44 are seen to be strongly reduced by smoothing the masking between the target pixels and the rest of the Fourier domain image, but they still largely remain.

in the end a much more complex solution is needed that can distinguish between the contribution to the frequencies associated with the background from contributions to the same frequencies that arise in the true astronomical image. In particular, the edges of stars, since they are better represented by a Gaussian than a box function, prove problematic as they cover a large range of frequencies in the Fourier space image, from the high frequencies of the point source star, to the lower values associated with the extended scattered light around the star in an image. A simple masking method, though possible, requires us to know the parameters of the background we are trying to account for in the first place, bringing us back to the original problem of tackling the scattered light using a direct fitting approach.

5 Production of Cleaned Hα Images and Mosaicking

To reach the goal of producing clean and smooth $H\alpha$ mosaics from the IPHAS data, the background model fitting method described in Section 3 (which was trialled and tested on i band data, free of diffuse nebulosity) must be adapted to work on $H\alpha$ images. Here, we present the work needed to achieve this goal, describing the hurdles faced and overcome. The overall algorithm arrived at is summarised in the flow chart in figure 46.

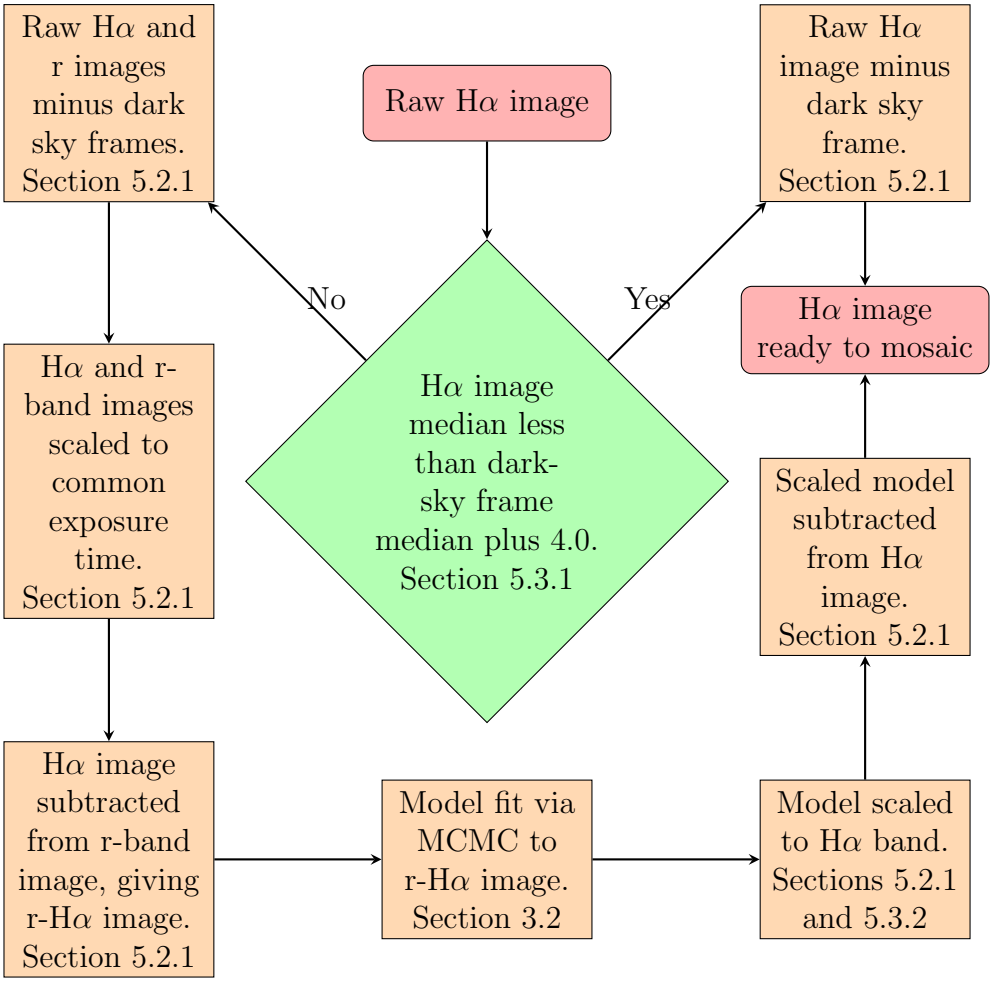

Figure 46: A flow chart detailing the methodology used in the final version of the background sky correction code.

5.1 How Hα images have been cleaned

We cannot directly fit a model to the Hα background because the diffuse nebulosity that we want to preserve in these data is indistinguishable from other sources of background in the re-binned mapping. While small angular scale nebulous features, like much of that found in supernova remnants like Simeis 147, won't affect the background modelling, there will be circumstances in which very extended diffuse Hα emission could contribute to the background gradient. Since the Hα filter falls entirely within the r band wavelength range (figure 7), the Hα (and [NII] 6548, 6584 emission) can be removed from any given r band image by simply subtracting from it its partner Hα image. A background model can then be constructed to fit the '$r - H\alpha$' result. This in turn can be scaled appropriately to yield a model for the background suitable for subtraction from the Hα image, under the assumption that the spatial variation across the four CCDs will be the same in both the r and Hα exposures taken one after the other.

5.2 The sources of non-stellar counts in r and Hα band images

To remove the unwanted components of diffuse light in an Hα image, accurately, we need to understand how the pixel values in a typical image are built up.

The total count, $r(x, y)$ in pixel (x, y) in an r band image is made up as follows:

$$\frac{r(x,y)}{t_r} = d(x,y)_r + c(x,y)_m + c(x,y)_{H\alpha} \tag{11}$$

where t_r is the exposure time, and the terms on the right hand side are (in order) the count rates for dark-sky, scattered moonlight, and astronomical line emission. In the interests of simplicity, we leave out of consideration the continuum starlight found in point sources. In the r band, t_r is either 10 or 30 seconds (after the start of 2004, the exposure time was raised to the higher number).

The total count, $H\alpha(x, y)$, in a narrowband image is described by the analogous relation:

$$\frac{H\alpha(x,y)}{t_{H\alpha}} = d(x,y)_{H\alpha} + \frac{1}{k}.c(x,y)_m + c(x,y)_{H\alpha} \tag{12}$$

where $t_{H\alpha}$ is always 120 sec. The constant k multiplying the scattered moonlight count rate is the factor needed to scale down the rate passing through the r filter to that appropriate to the narrowband. In principle this can be

computed from a knowledge of the spectrum of scattered moonlight and the filter bandpasses. Treating scattered moonlight as having the same spectrum as direct moonlight (Krisciunas and Schaefer, 1991) (as given by Jones et al. (2013)), the value of this constant is 12.84 (Cramer et al., 2013)- a number not very different from the ratio of the filter FWHM values.

Typical dark sky count rates ($d(x, y)$) in the two filters can be determined from merging r and Hα images, free of nebular emission, that were obtained on moonless nights. The goal here is to recover $c(x, y)_{H\alpha}$. In the two equations above, this amounts to one unknown while the other is the scattered moonlight 'image', $c(x, y)_m$. This is a system that can be solved, given that all other quantities in the equations are known or knowable.

The method used to clean the Hα images is a series of image manipulations incorporating background fitting of the kind described in chapter 3.2 in order to reduce noise in the final outcome. The next section traces this in more detail.

5.2.1 Cleaning r and Ha images and removing nebulosity

Before any further data processing takes place, the images in r and Hα : first are confidence cleaned as in Section 3.1.2; masked for bright stars as described in Section 3.1.3; have their dark sky templates subtracted as in Section 3.1.1. This last step of removing dark sky accounts for $d_r(x, y)$ and $d_{H\alpha}(x, y)$ in equations 11 and 12.

It is not possible to fit a model directly to the background of images in r or Hα because of the nebulous structure represented by $c_{H\alpha}$. This can be overcome by removal of the Hα pixel counts from their partner r-band image. Before subtracting one from the other, both the r and Hα images must be scaled to the same exposure time. To simplify things, all images are scaled to 1s exposures. Once both an r-band and an Ha image are scaled to the same exposure time, the Hα image is subtracted from an r-band image which leaves and r-Hα residual image where:

$$\frac{r(x, y)}{t_r} - d_r(x, y) - \frac{H\alpha(x, y)}{t_{H\alpha}} + d_{H\alpha}(x, y) = (1 - \frac{1}{k})c(x, y)_m \qquad (13)$$

A model is then fit to this r-Hα residual image in the same way as in Section 3.2. The factor of $1 - \frac{1}{k}$ accounts for the r-band moonlight count rate lost thanks to the subtraction of the narrowband image from the r-band image. A correction for this can be made to the model fit of the r-Hα residual image by dividing by $1 - \frac{1}{k}$. The model is scaled back to the narrow-band filter width in this way, before being scaled up to an equivalent 120sec Hα

exposure time. At this point, the model is subtracted straight off its partner (confidence cleaned and dark-sky subtracted) Hα image. The result should have a flat background at a counts level close to zero, retaining all of its original stellar counts and nebulous features, ready for mosaicking.

5.3 Issues faced

When applying the method in Section 5.1 to the 580 IPHAS images around S147, initial results came out with a range of residual background levels after model subtraction, as shown in figure 47. There are two key features in this plot that point to issues with the model fitting process.

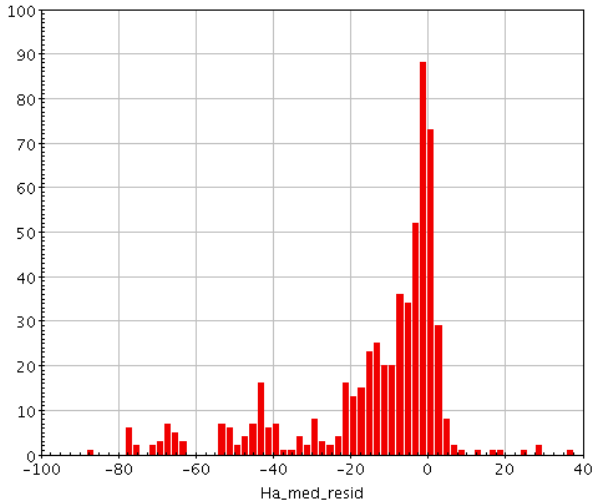

Figure 47: The median of the residual counts for each of the 580 background-corrected Hα image in the S147 region at first attempt. The spread is large, and there is a skew towards negative values.

5.3.1 Negative skew in the corrected Hα residual counts

The first issue faced, the negative skew of the residual medians, is considered here, along with its cause and a fix that avoids over-processing already clean data unnecessarily.

When the dark sky templates are subtracted from frames obtained in dark time, there can be over-subtraction. If this happens, a substantial number of pixel count rates become negative. Any negative counts created this way in an Hα image turn into added counts at the point where the Hα image is subtracted from its matching r-band image. Therefore, when the background model is fit to the r-Ha image, it is set too high. At the next step, when

subtracting this model from the Hα image, a significant over-subtraction of the background can occur. The median of the residual background in such cases is then negative. This seems to be behind the negative skew in figure 47.

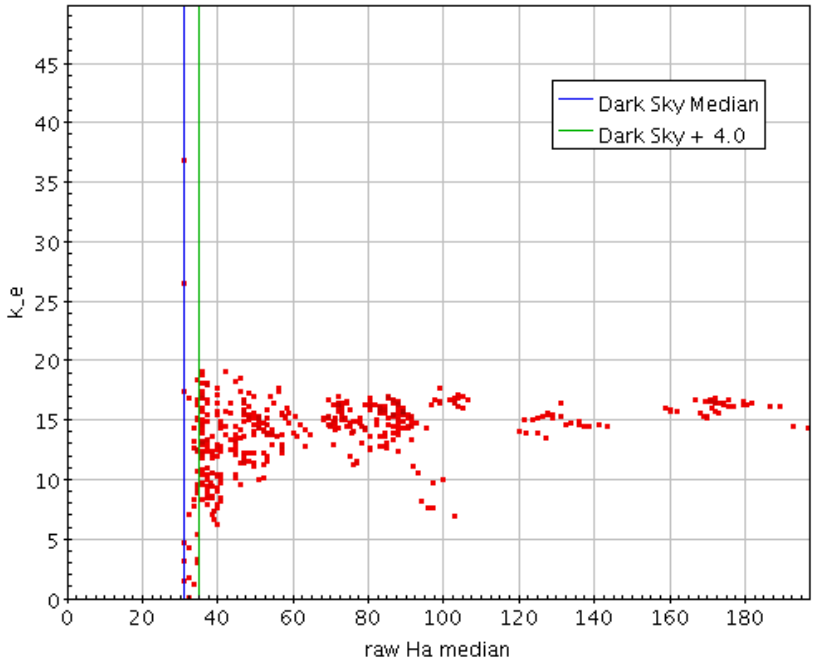

Figure 48: The empirical scale factor, k_e (see text), for each r and Hα pair of images plotted against the raw Hα image median count for each of the 580 S147 images. In cases where the raw median is approximately the same as the median of the dark sky templates, k_e becomes unreliable as the remaining counts after dark-frame subtraction are close to zero. The blue line shows the median dark-sky value and the green line represents the cut-off background count value of 4.0 counts above the dark-sky median.

To overcome the negative skew in the final corrected background levels, any raw Hα images that have an Hα median background (or sky) value within 4.0 counts of the median of the relevant dark-sky template, the cleaning process ends after dark-sky subtraction, and there is no background model fitting step. A limit of 4.0 is used since cases where the median of the raw image is within 4.0 counts of the median of the dark-sky template begin to yield anomalous empirical estimates for k, as seen in figure 48. We remind the reader that k is the expected ratio between the r-band and Hα band moonlight count rates. Its empirical counterpart, k_e, is estimated as the ratio between the median count rates of the dark-sky subtracted r and Hα frames. The spread in empirical scale factor seen for images where the Hα

median background value is within 4.0 counts of the dark-sky median (31, in this case) in figure 48 is due to the residual counts in r and Hα both approaching zero. In the 47 such cases in this example, this signals that there is little or no moonlight to remove.

By stopping the cleaning process when the sky background is evidently dark, we avoid any unnecessary processing. In such cases, the dark-sky frame is scaled to the raw image by using the ratio of the background median pixel count to the corresponding dark-sky frame median pixel count. This scaling is necessary to ensure the final image in r and Hα has a zero background since, as described in Section 2.1, the airglow in these frames can vary by up to $\sim 10\%$. In figure 48, images with k_e outside the range of 0 to 20 are all cases that have been over-processed unnecessarily.

5.3.2 The scaling, k, of r-band to narrowband moonlight levels

The second issue to be faced was the difference between the expected ratio, k, and the empirical values, k_e. The distribution of empirical values is presented in figure 49.

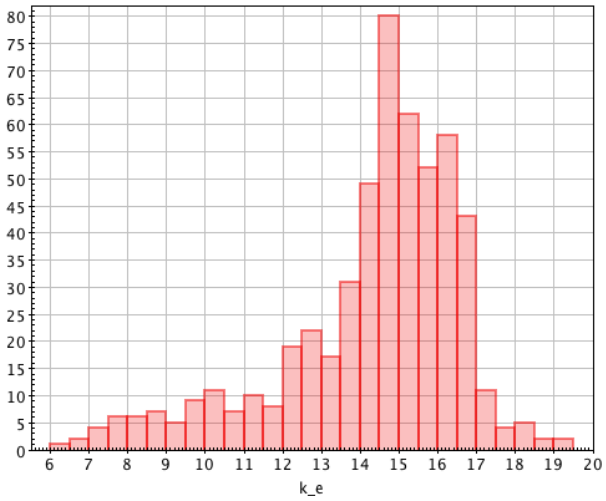

Figure 49: The empirical scale factor distribution for 580 images in the S147 region. The expected value is 12.84, while the empirical data favour ∼15.

This range in k_e also contributes to the spread of residuals seen earlier in figure 47. The scalefactor used to convert the r band moonlight levels to Hα levels varies because the direct moon spectrum (that we used to determine $k = 12.84$) is not identical with the scattered moon spectrum, and because in unsteady sky conditions the images will not sit in a constant ratio, even though they are observed consecutively. More often than not $k = 12.84$

is an underestimate causing the subtraction of the background model off Hα to overshoot. Rayleigh scattering in the atmosphere is more efficient at shorter/bluer wavelengths, implying that a scattered moonlight background would be expected to be bluer and that there will be some variation in the SED change according to the geometry of the scattering (see Jones et al. (2013)). As the narrow-band filter sits toward the red end of the r bandpass, this effect would indeed cause the value of k to rise, rather than fall. The scale of effect seen is in the region of a 10-15 percent increase. Unfortunately there is not a library of flux-calibrated scattered moonlight spectra that we can access to confirm the validity of this interpretation. To deal with this we use k_e in the algorithm instead.

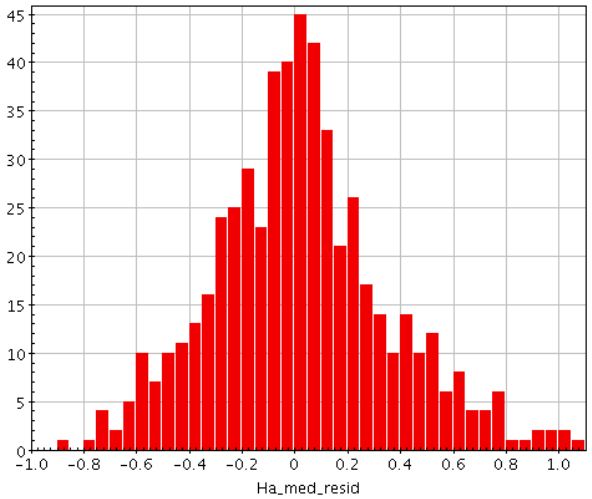

Figure 50: The median residual count in each of the 580 S147 region images after cleaning, implementing both scaling for the dark-time frames, and using the empirical moonlight scaling, k_e for each image. There is huge improvement from figure 47, in which there is a large sprread of values and a negative skew. Now, the residuals peak at 0 by design, and vary very little (± 1 count).

Implementing an empirical k_e, and scaling the dark-sky frames while avoiding over processing clean data, the residuals of the resulting data in the S147 region come out as in figure 50. These residuals now vary by tiny amounts, and are within the range of -1 and 1 without skew. To gauge the success of this procedure, we turn to appraisal of the cleaned and mosaicked images in the S147 region.

5.4 Montage

Throughout this project, we use MONTAGE to mosaic IPHAS images together into a complete mosaic, re-projecting the images onto a common coordinate system and adjusting the background levels to produce a smooth image. MONTAGE is a piece of freely available software that allows control over a number of parameters throughout the mosaicking process. The steps in which to use MONTAGE to go from a collection of images to a complete mosaic are as follows:

Figure 51: An example mosaic in Hα using IPHAS data of the Rosette nebulae. North is up and East is left, covering $2^\circ \times 2^\circ$.

- Step 1 - Creates a list of the images that require mosaicking.
- Step 2 - Creates the FITS header template file for mosaic using the image list created in step 1. At this point it is possible to specify to MONTAGE to align all images with North up. The header template defines the final mosaic in terms of its projection method, images size, orientation, and pixel scale. The header template can be created by hand to produce an image with a desired specification, or it can be created automatically by MONTAGE.
- Step 3 - Each image is then re-projected to the size and scale defined in the header template. As well as producing the re-projected image,

this step creates an 'area' image which consists of the fraction input pixel sky area that went into each output pixel.

- Step 4 - These files are then used to create an image metadata table which contains the geometry information for each image being mosaicked and is used in several other steps. At this point a FITS mosaic becomes an optional output as a means to check the progress of the process and that no initial errors have occurred. These images, however, are created before any internal background correction has been done, and so they provide guidance only.

- Step 5 - MONTAGE compiles a list of all the overlapping pairs of images by comparing the area of each image to every other image.

- Step 6 - The 'difference' between the overlapping regions is calculated and stored in individual images, one for each pair of overlapping images.

- Step 7 - MONTAGE uses the difference images produced in the previous step to fit planes for each image that will allow it to smooth out the overlapping areas of the mosaic. These are stored in a table. It can be specified here as to whether MONTAGE is to match the slopes between two overlapping images, or whether it should only calculate level adjustments between the two.

- Step 8 - Uses the table of difference parameters from step 7 to determine the set of corrections that need to be applied to the background of each image to achieve the smoothest fit across the final mosaic.

- Step 9 - These background corrections are then applied and a new set of (corrected) images is produced.

- Step 10 - With these final, corrected images produced, the image metadata file created in step 4 is adjusted to remove any images that didn't make it to this point (typically only 1 or 2, due to not having any overlap).

- Step 11 - The images are then co-added to form the final output FITS mosaic, with an example mosaic in figure 51.

5.5 Zero Point correction

The final step before mosaicking the cleaned data is to adjust for their zero points (ZPs). The ZP is the magnitude of an object in an image that will

produce one count per second such that the magnitude of an arbitrary object in the image is:

$$mag = -2.5 \log(pixelvalue) + ZP \tag{14}$$

The light has to pass through the entire atmosphere before forming the image, and so is subject to transparency variations due to clouds, and extinction due to dust and aerosols. The pipeline processing of the IPHAS data uses standard fields to make a routine correction for airmass and provides a ZP that, in principle, accounts for these variations. For those exposures that were adopted in IPHAS DR2, the zero point has been calibrated as described in Barentsen et al. (2014). Fields obtained more recently have so far only been calibrated on a nightly basis, within the CASU pipeline.

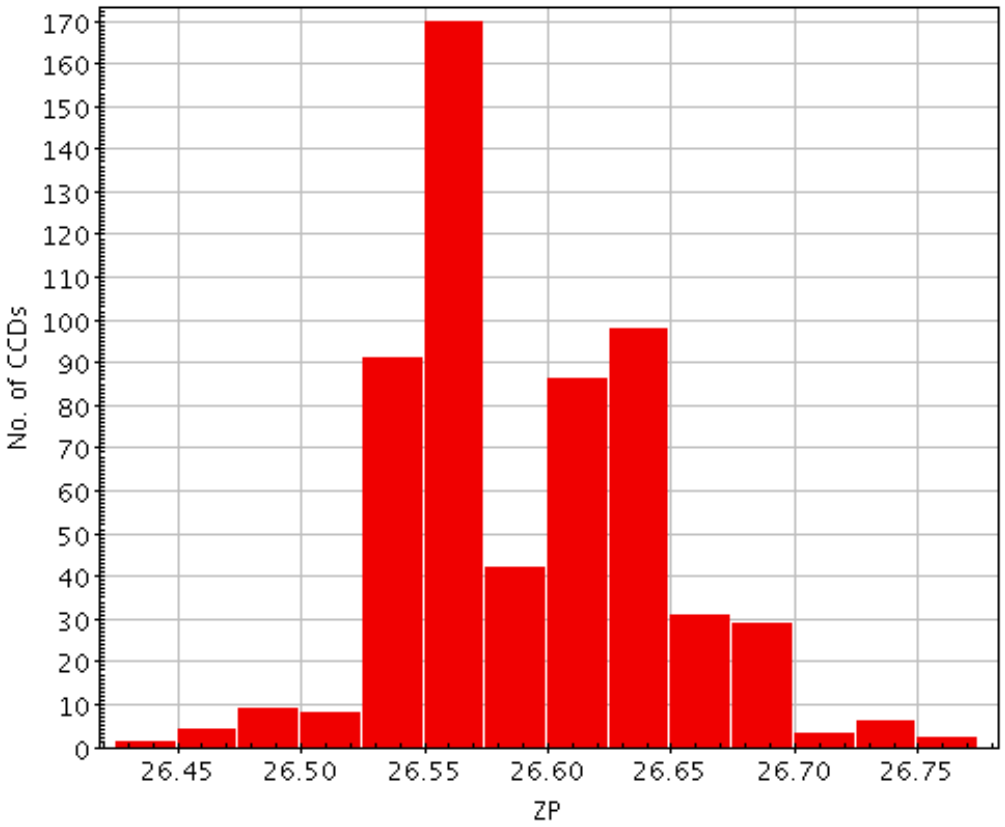

Figure 52: The distribution of CCD zero points taking up the $4° \times 4°$ *area around Simeis 147.*

Once the cleaning process is complete, what will remain will be these transparency variations that the ZPs track. It is important to bring the

images onto a common ZP before mosaicking because, since magnitudes are logarithmic, a change of e.g. 0.75 in ZP across two images would correspond to a factor of 2 change in counts according to:

$$SF_{ZP} = 10^{-0.4(ZP-ZP_{ref})} \tag{15}$$

where SF_{ZP} is the scale factor bringing the image onto the common zero point ZP_{ref}. ZP_{ref} is taken as the median of the ZPs in the region from figure 52.

We use equation 15 to bring the counts in all images on to the same common zero point of 26.58, which corresponds to an equivalent ZP for a 1sec exposure of 21.38.

In some of the worst cases of bad data in the catalogue, replacement observations (as mentioned in section 1.5.3) have been collected in December 2015 which have not yet been subject to global calibration. These images share a zero point of 21.65.

5.6 Final mosaic

The final version of the mosaic of S147 is displayed in figure 53. The image covers $4^\circ \times 4^\circ$ on the sky. Before mosaicking, the images were binned by a factor of 4, reducing the processing time. A comparison in figure 54 shows the extent of the cleaning that has taken place, with bright scattered moonlight in the original data causing much of the structure of S147 to no longer be visible.

Figures 55 and 56 shows four zoomed-in sections of figure 53, illustrating the level of detail visible once the scattered light is removed. Images like this can give great insight into the morphology of SNRs and the history of their evolution. The data here is binned by a factor of 4, so the potential of the un-binned data is even greater. Lens affects can be seen around the saturated brightest stars - unfortunately these are almost impossible to remove.

Finally, in figure 57 we see a before-and-after case where the same cleaning method is applied to another patch of sky. In this case, the Heart nebula (NGC 896) shows the ability of the cleaning process to be applied across the sky. The most notable improvements in this image is the appearance of structure in the south, which appears without having to over-stretch the image and saturate the brightest structure in the north.

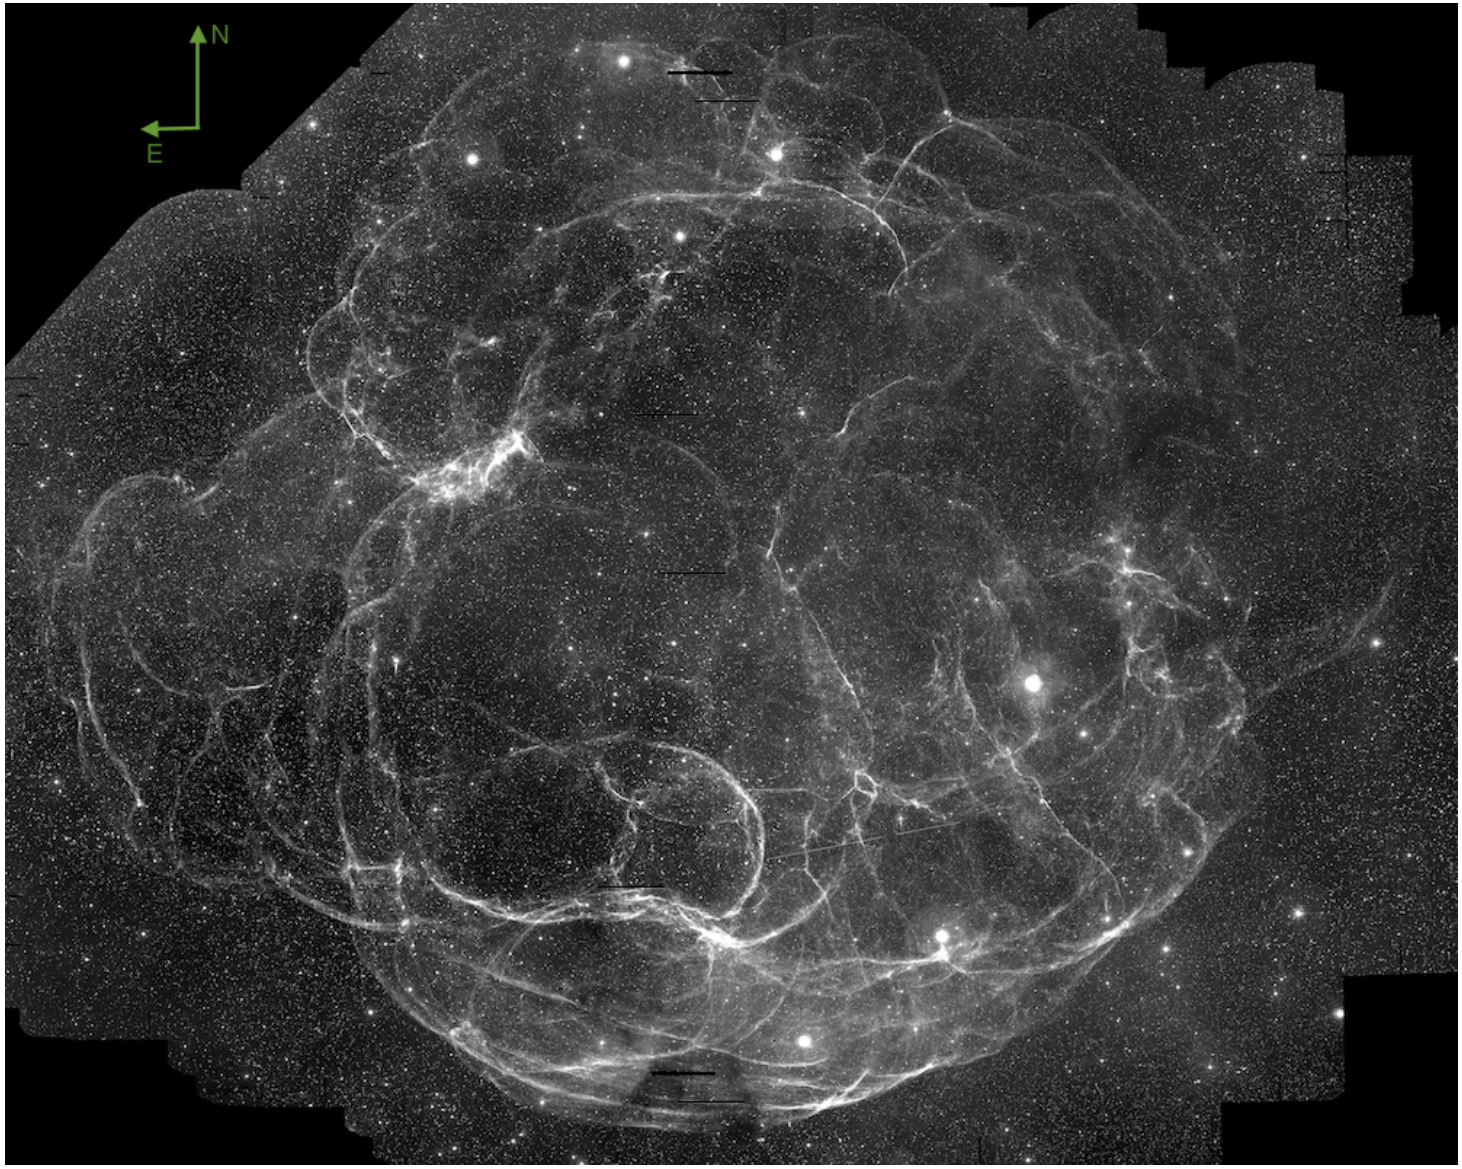

Figure 53: The final mosaic of S147, after all data has been through the cleaning process, replacement data has been added, and the images have been moved on to a common zero point. The image covers $4^\circ \times 4^\circ$, North is up and East is left.

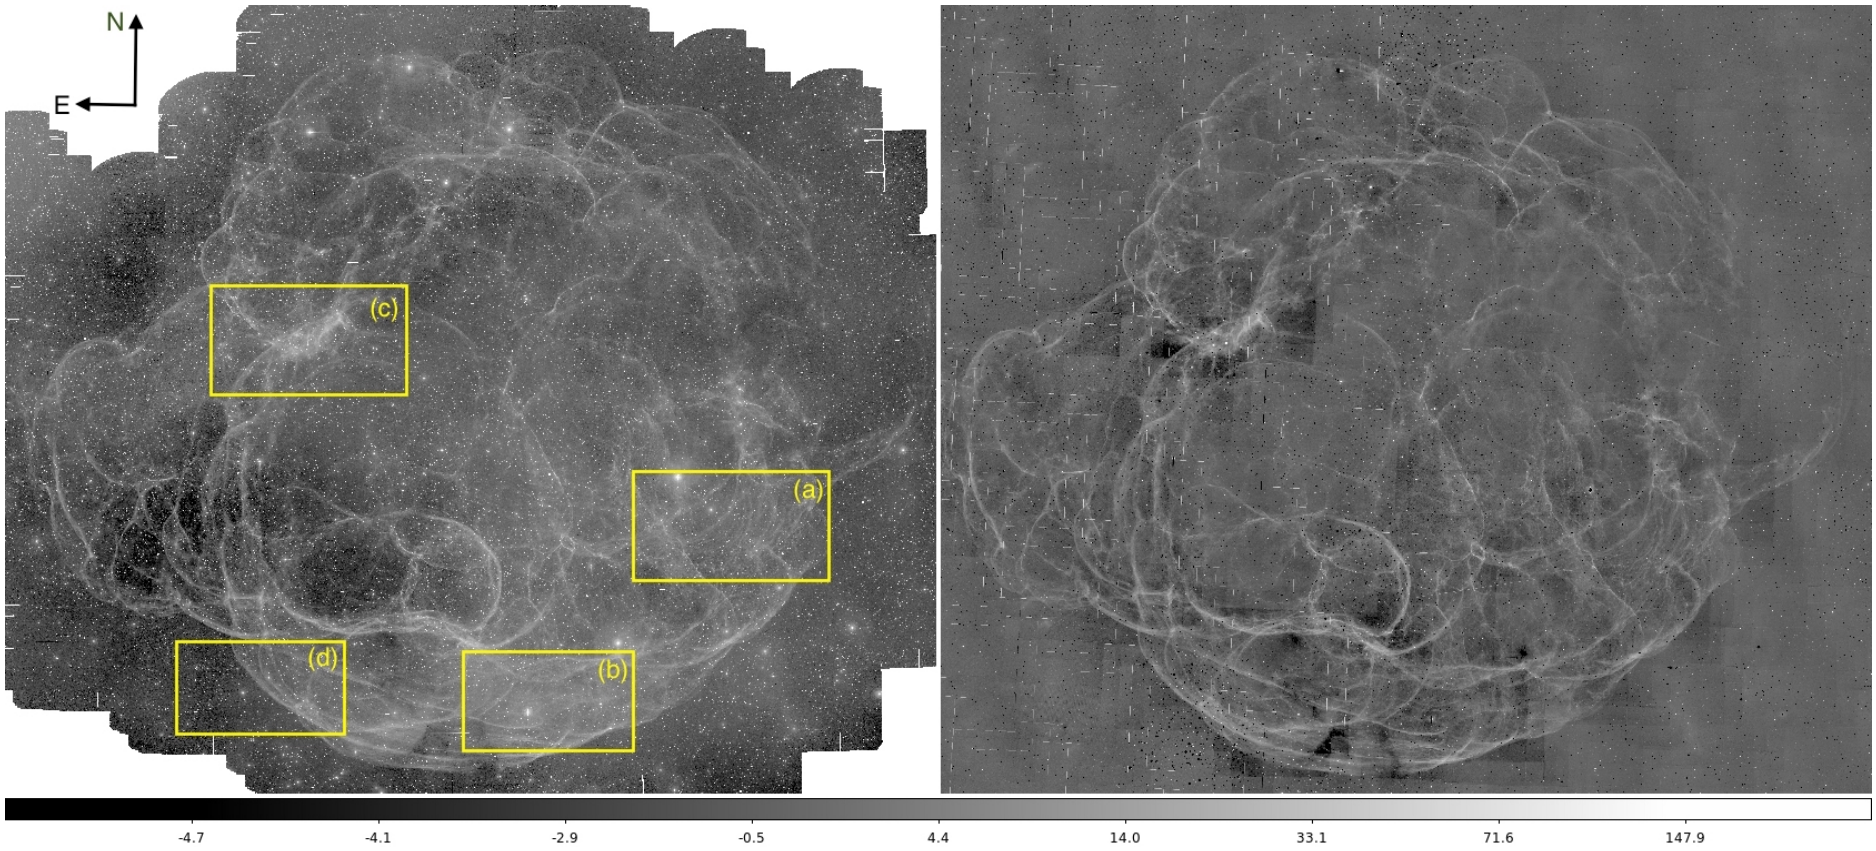

Figure 54: S147 after (left) and before (right) cleaning. The images cover $4° \times 4°$, North is up and East is left. Both images are scaled to the same colour bar. The right-hand image has been r-band continuum subtracted as the first attempt to produce mosaics from high moon illumination areas during the early days of IPHAS. The left hand image does not suffer from the poor smoothing of the background seen in the right-hand image, and does not suffer from losses of stars and diffuse material that comes with continuum subtraction.

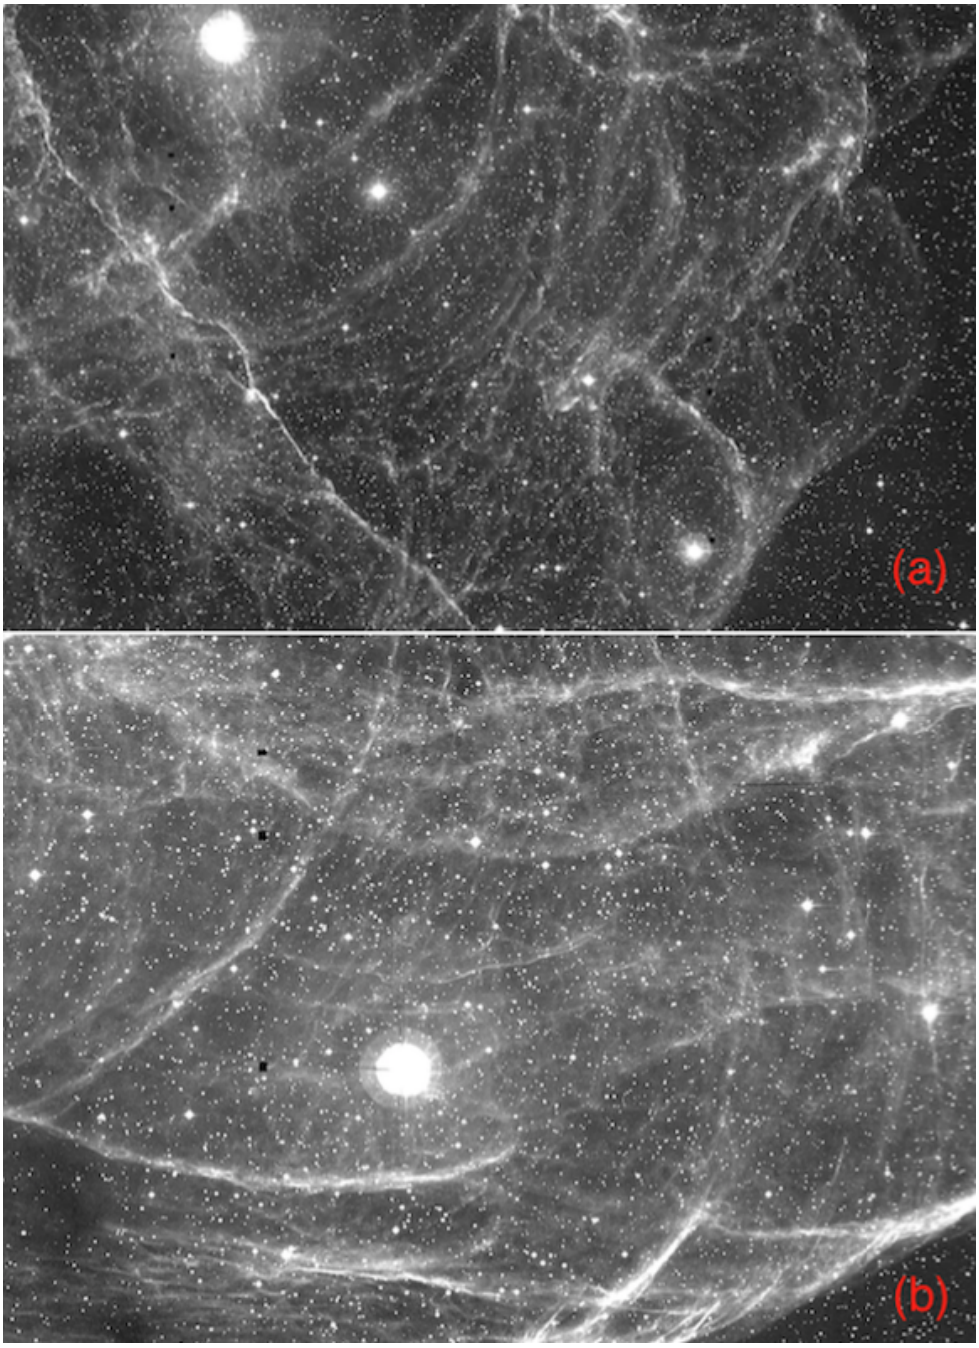

Figure 55: Two sections of the final mosaic (figure 53), with their locations marked in figure 54, to highlight the detail available in the cleaned images. North is up and East is left in all.

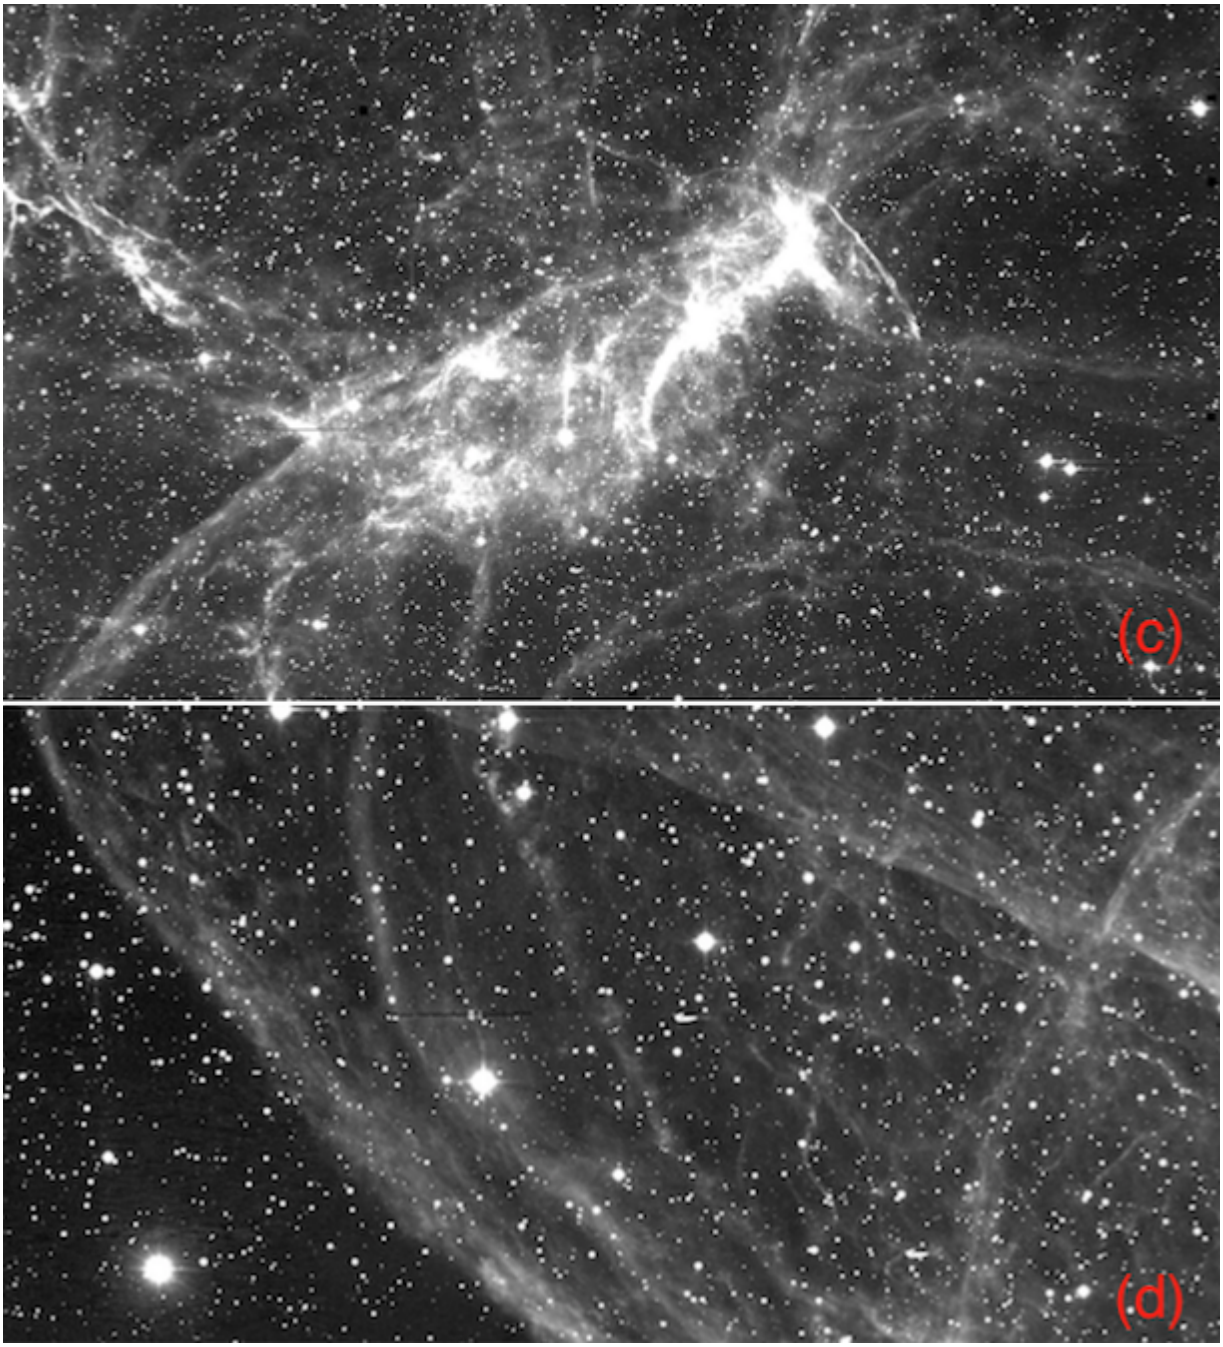

Figure 56: Two sections of the final mosaic (figure 53), with their locations marked in figure 54, to highlight the detail available in the cleaned images. North is up and East is left in all.

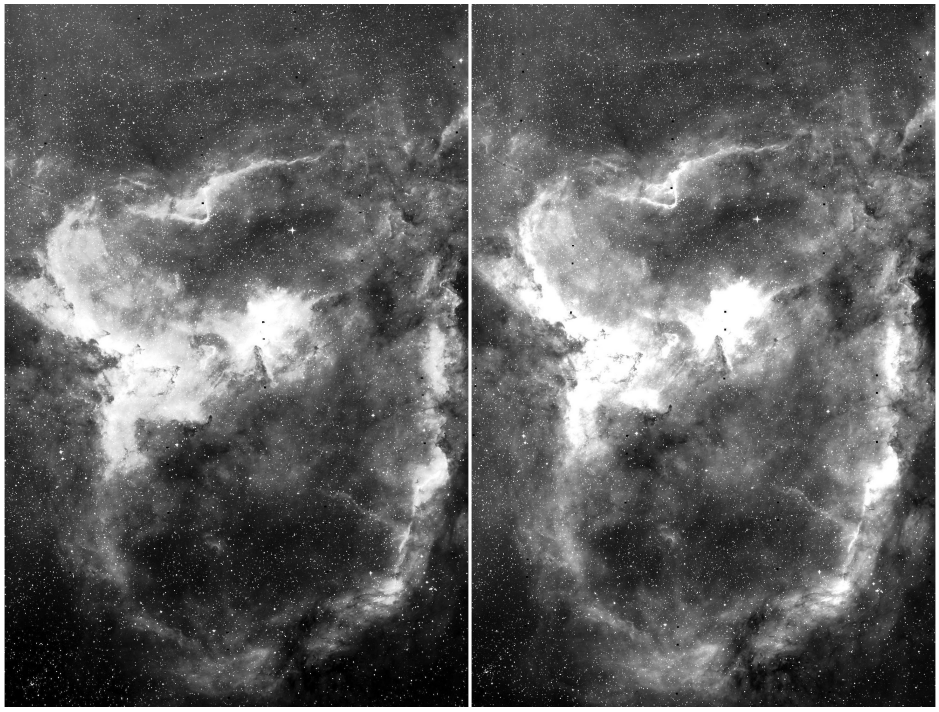

Figure 57: (Left) the Heart nebula without cleaning and (right) after passing through the cleaning process. North is up, east is left. Though the quality of the raw data in this region is generally good, cleaning has made it possible to see structure in the south of the region that was otherwise invisible when scattered moonlight was present in the mosaic. The greyscales are the same for both images. The images cover $1.5^\circ \times 2^\circ$*, North is up and East is left.*

6 Conclusions

The study of Hα allows us to better understand the formation, evolution and morphology of many structures and energetic processes within the ISM. Crucial to the scientific study of wide-area Hα emission are large-scale, high-resolution mosaics clean of all unwanted sources of non-astronomical light.

We set out with the goal of being able to produce large-scale mosaics of IPHAS data cleaned of unwanted contributions to the background, particularly the large-scale gradients introduced by scattered moonlight. The final algorithm adopted uses Hα and r-band frames to achieve this. The first step in the scheme was to create dark-time image templates that could be subtracted from the images to remove the contribution of dark-sky sources, such as airglow. The Hα data is subtracted from the r-band to remove nebulosity from the broadband which permitted a model to be fit to the r background without influence from large-scale Hα structure. The brightest stars are also cut out of the images, before a linear, 2-dimensional model is fit to the remaining background via Markov Chain Monte Carlo fitting. This model, originally tested on the i-band data that is free of nebulosity, is scaled to the Hα filter width first using a scale factor of 1/12.84, originating from the relative ratio of the predicted in-band fluxes due to moonlight. This scale factor is replaced by an empirical scale factor, k_e (section 5.3.1) as this was found to give results of the desired quality. Next, the scaled model is subtracted out of the Hα data leaving us with a cleaned image. Finally, the cleaned images are shifted onto a common zero point to account for transparency variations, before finally being mosaicked together.

We have shown that the cleaning method incorporating 2-D background fits, greatly improves the quality of the outcome: the RMS values describing the quality of the fits found by the MCMC code fall from 3-4 to less than 1. Later in the process we find the median of the residual background in the cleaned images generally lies between -1 to 1, to be compared with the peak counts of $\sim$30 in the case of S147. The final assessments came in the form of visual analysis of the final mosaic, looking for regions that were not completely flat and comparing the features now visible in the data with that before the cleaning process. The residual backgrounds in the resulting images - and therefore mosaics - are approximately flat. Mosaicking the data together in the region around the supernova remnant S147 shows the ability of this method to produce large-scale and clean mosaics, without being influenced by bright stars or nebulosity. The method is also seen to work well when applied to other regions without any specific preparation. Our method allows us to retain the smallest detail in the data and, since we avoid removing any nebulosity, we expose and resolve some of the faintest

and most intricate morphology seen from the ground.

Perhaps the key weakness of the process as applied so far is seen in images that are awash with nebulosity. In such cases, the empirical scale factor that converts between r and Hα counts is likely to be smaller than it should be, resulting in an 'over-subtraction' when it is removed from the Hα data. Since S147 is a late-stage SNR, and therefore is made up of dense, thin filamentary structure, this doesn't pose an issue to our example area. A possible fix for this problem could be to set a lower limit on k_e of 12.84, on the basis that ratios below this would be implausibly red for scattered moonlight.

References

S. B. Anderson, B. J. Cadwell, B. A. Jacoby, A. Wolszczan, R. S. Foster, and M. Kramer. A 143 Millisecond Radio Pulsar in the Supernova Remnant S147. , 468:L55, September 1996. doi: 10.1086/310218.

G. Barentsen et al. The second data release of the INT Photometric Hα Survey of the Northern Galactic Plane (IPHAS DR2). , 444:3230–3257, November 2014. doi: 10.1093/mnras/stu1651.

C. Benn and S. Ellison. La Palma Night-Sky Brightness. April 2007.

C. Bertout and J. Bouvier. T Tauri disk models. In B. Reipurth, editor, *European Southern Observatory Conference and Workshop Proceedings*, volume 33 of *European Southern Observatory Conference and Workshop Proceedings*, pages 215–232, September 1989.

R. N. Bracewell. *The Fourier Transform and its Applications*. McGraw Hill, November 1999.

E. Oran Brigham. *The Fast Fourier Transform and its Applications*. Prentice Hall, 1988.

C. E. Cramer et al. Precise measurement of lunar spectral irradiance at visible wavelengths. *Journal of Research of the National Institute of Standards and Technology*, 118:396–402, 2013.

B. Dinçel, R. Neuhäuser, S. K. Yerli, A. Ankay, N. Tetzlaff, G. Torres, and M. Mugrauer. Discovery of an OB runaway star inside SNR S147. , 448: 3196–3205, April 2015. doi: 10.1093/mnras/stv124.

B. T. Draine. *Physics of the Interstellar and Intergalactic Medium*. Princeton, 2011.

J. E. Drew et al. The INT Photometric Hα Survey of the Northern Galactic Plane (IPHAS). , 362:753–776, September 2005. doi: 10.1111/j.1365-2966.2005.09330.x.

G. B. Field, D. W. Goldsmith, and H. J. Habing. Cosmic-Ray Heating of the Interstellar Gas. , 155:L149, March 1969. doi: 10.1086/180324.

D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. , 125:306–312, March 2013. doi: 10.1086/670067.

E. Fuerst and W. Reich. Multifrequency radio observations of S 147. , 163: 185–193, July 1986.

E. González-Solares et al. Data Mining Large Surveys: The IPHAS Early Data Release. In R. W. Argyle, P. S. Bunclark, and J. R. Lewis, editors, *Astronomical Data Analysis Software and Systems XVII*, volume 394 of *Astronomical Society of the Pacific Conference Series*, page 197, August 2008.

V. V. Gvaramadze. PSR J0538+2817 As The Remnant Of The First Supernova Explosion in a Massive Binary. In *IAU Joint Discussion*, volume 2 of *IAU Joint Discussion*, August 2006.

F. J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. 1978.

E. Høg, C. Fabricius, V. V. Makarov, S. Urban, T. Corbin, G. Wycoff, U. Bastian, P. Schwekendiek, and A. Wicenec. The Tycho-2 catalogue of the 2.5 million brightest stars. , 355:L27–L30, March 2000.

M. Irwin and J. Lewis. INT WFS pipeline processing. , 45:105–110, January 2001. doi: 10.1016/S1387-6473(00)00138-X.

A. Jones, S. Noll, W. Kausch, C. Szyszka, and S. Kimeswenger. An advanced scattered moonlight model for Cerro Paranal. , 560:A91, December 2013. doi: 10.1051/0004-6361/201322433.

M. Kramer, A. G. Lyne, G. Hobbs, O. Löhmer, P. Carr, C. Jordan, and A. Wolszczan. The Proper Motion, Age, and Initial Spin Period of PSR J0538+2817 in S147. , 593:L31–L34, August 2003. doi: 10.1086/378082.

K. Krisciunas and B. E. Schaefer. A model of the brightness of moonlight. , 103:1033–1039, September 1991. doi: 10.1086/132921.

T. A. Lozinskaya. Optical observations of supernova remnants: the filamentary nebula Simeiz 147. , 20:19, February 1976.

T. A. Lozinskaya. *Supernovae and Stellar Wind in the Interstellar Medium*. American Institute of Physics, 1992.

C. F. McKee and J. P. Ostriker. A theory of the interstellar medium - Three components regulated by supernova explosions in an inhomogeneous substrate. , 218:148–169, November 1977. doi: 10.1086/155667.

R. G. McMahon, N. A. Walton, M. J. Irwin, J. R. Lewis, P. S. Bunclark, and D. H. Jones. The INT wide field imaging survey (WFS). , 45:97–104, January 2001. doi: 10.1016/S1387-6473(00)00137-8.

A. B. Meinel, II. OH Emission Bands in the Spectrum of the Night Sky. II. , 112:120, July 1950a. doi: 10.1086/145321.

I. A. B. Meinel. OH Emission Bands in the Spectrum of the Night Sky. , 111:555, May 1950b. doi: 10.1086/145296.

D. Osterbrock. Introduction to the interstellar medium. *UCSC course notes*, 1984.

A. P. Phillips and P. M. Gondhalekar. Further ultraviolet observations of interstellar gas associated with the supernova remnant S 147. , 202:483–493, February 1983. doi: 10.1093/mnras/202.2.483.

A. P. Phillips, P. M. Gondhalekar, and J. C. Blades. High-velocity gas associated with the supernova remnant S147. , 195:485–495, May 1981. doi: 10.1093/mnras/195.3.485.

J. M. Pittard. Self-sealing shells: blowouts and blisters on the surfaces of leaky wind-blown bubbles and supernova remnants. *Monthly Notices of the Royal Astronomical Society*, 435(4):3600–3613, 2013. doi: 10.1093/mnras/stt1552. URL `http://mnras.oxfordjournals.org/content/435/4/3600.abstract`.

S. P. Reynolds. Supernova Remnants at High Energy. , 46:89–126, September 2008. doi: 10.1146/annurev.astro.46.060407.145237.

E. L. Robinson. The structure of cataclysmic variables. , 14:119–142, 1976. doi: 10.1146/annurev.aa.14.090176.001003.

L. Sabin et al. New Galactic supernova remnants discovered with IPHAS. , 431:279–291, May 2013. doi: 10.1093/mnras/stt160.

J. L. Sauvageot, J. Ballet, and R. Rothenflug. EXOSAT observation toward S147 and A0535 + 26. , 227:183–187, January 1990.

J. Silk and G. Wallerstein. High-velocity gas in supernova remnants. II. Shajn 147. , 181:799–804, May 1973. doi: 10.1086/152090.

S. J. Smartt. Progenitors of Core-Collapse Supernovae. , 47:63–106, September 2009. doi: 10.1146/annurev-astro-082708-101737.

Y. Sofue, E. Furst, and W. Hirth. Radio Continuum Observations at 5-GHZ of the Supernova Remnant S147. , 32:1, 1980.

L. B. F. M. Waters. Herbig Ae/Be Stars. In M. Kraus and A. S. Miroshnichenko, editors, *Stars with the B[e] Phenomenon*, volume 355 of *Astronomical Society of the Pacific Conference Series*, page 87, December 2006.

K. W. Weiler and R. A. Sramek. Supernovae and supernova remnants. , 26: 295–341, 1988. doi: 10.1146/annurev.aa.26.090188.001455.

This research made use of MONTAGE. It is funded by the National Science Foundation under Grant Number ACI-1440620, and was previously funded by the National Aeronautics and Space Administration's Earth Science Technology Office, Computation Technologies Project, under Cooperative Agreement Number NCC5-626 between NASA and the California Institute of Technology.