Auroral Image Data Analysis

Björn Gustavsson

September 29, 2011

Auroral Image Data Analysis
Typeset by the author in LaTeX.

ISSN TBD:1.42413562
ISBN TBD:2.718281828
Printed at the Printing Press
Adress, Adress, Country
September 29, 2011

Contents

List of Figures

Chapter 1

Introduction

1.1 Image analysis; a panoramic overview

An image can contain a lot of information. To fully understand and utilise the possibilities in image analysis it is necessary to have a good grasp of a variety of aspects of the imaging process. To be able to filter image data it is necessary to understand the noise characteristics of different steps in imaging, to be able to do stereoscopy and triangulation from images it is necessary not only to understand the geometric aspects of optics and imaging but one also need to know how to estimate the geometric properties of the imaging with high enough accuracy. To be able to correctly compare different image intensities in images taken with different filters using different cameras one need not only know how to calibrate the absolute camera sensitivity but also how the image sensitivity varies over the image plate.

These and other topics will be addressed in the following chapters.

The presentation here use a different mathematical model for imaging compared to what I have found in other texts on image processing. What this lacks in compactness it gains in being more general and its wider applicability. As in all fields of physics it is as important to understand how far it is possible to stretch the use of ones models, and in what way they break down when stretched to (or, o horror, beyond) their limits.

The reason I wrote this is that I could not find the information I needed in the very good books I read on optics, image processing, machine vision and computer graphics when I started my PhD studies in Auroral Physics - where my task was to develop data analysis methods for the Auroral Large Imaging System (ALIS) [*Brändström*, 2003]. This included everything from camera calibrations to tomographic estimates of the three-dimensional distribution of auroral volume emission rates from a limited number of view-points. So, this is not an attempt to replace or rewrite a standard text book on image processing like "Digital Image Processing" [*Gonzalez and Woods*, 1993], since Gonzalez and Woods already have written that one and have a few editions head start. Other excellent books I had much use for are "Image Acquisition" by *Burke* [1996] covering much of the optical/technical side of imaging, "Graphics Gems" by *Glassner* [1990] is a brilliant relief for the tired brain,

Figure 1.1: A famous example of the consequences of stretching models beyond their limits: The Tacoma Narrows Bridge collapsing because of driven oscillations in moderate winds. Also a stretch of association.

there everything of tricky geometry can be easily found and a more in depth cover of fundamental optics can be found in for example "Optics" by *Hecht and Zajac* [1987]. What I tried to do is to fill a small gap somewhere in a corner between the fields covered by the books mentioned.

Chapter 2

Camera models and image formation

Pin hole cameras and other black boxes.

2.1 Introduction

To fully understand the possibilities of image analysis it is necessary to understand the imaging process. In this chapter we will cover some fundamentals of the geometric properties of imaging, including a short repetition of coordinate transformations needed for calculating and interpreting images from a camera in arbitrary position and rotation. This is necessary to be able to determine in what direction an imaged

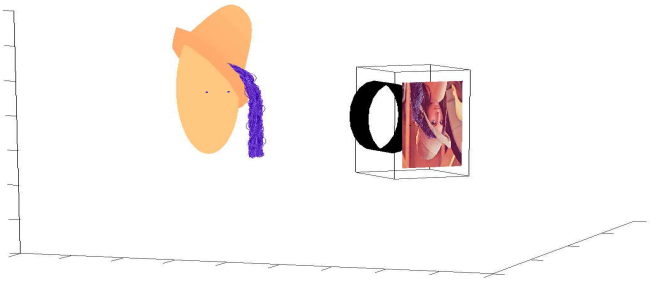

Figure 2.1: The geometry of imaging is imperative to all forms of spatial image analysis. Both for modelling of the image given a scene and for determining positions of objects in images.

object is, and when combining images from several view-points the position (from 2 or more view-points by stereoscopy) or internal structure (from many view-points by tomographic inversion). The main part is an in depth description of the imaging process, covering a description of a general camera model and its characteristics - that is a step by step description of how a three dimensional scene are projected down to an image plane.

Figure 2.2: It has been argued that the intricate geometry of the chandelier in the Arnolfini portrait by vanEyck must have been outlined by help of a camera obscura. Here we rather point out the accurate image in the curved mirror on the far wall.

2.2 Pre-history of imaging

> *Who would believe that so small a space could contain the image of all the universe? O mighty process! What talent can avail to penetrate a nature such as these? What tongue will it be that can unfold so great a wonder? Verily, none! This it is that guides the human discourse to the considering of divine things. Here the figures, here the colours, here all the images of every part of the universe are contracted to a point. O what a point is so marvellous! – Leonardo Da Vinci*

The simplest imaging device is the pinhole camera. It consists of a box with a small hole in one wall. Light entering the hole from the outside form an inverted image of the outside scene on the wall opposite the hole. The principle of image formation by this pinhole set-up have been known since at least 500 BCE when Mo Ti (or "Mo Tsu") described it.

Its potential as a drawing aid may have been familiar to artists by as early as the 15th century when its light gathering was improved by replacing the pinhole with a lens in a larger aperture. Leonardo Da Vinci described the camera obscura and it has been speculated that Johannes Vermeer made use of one, but the extent of their use by artists at this period remains a matter of debate. From what little I know

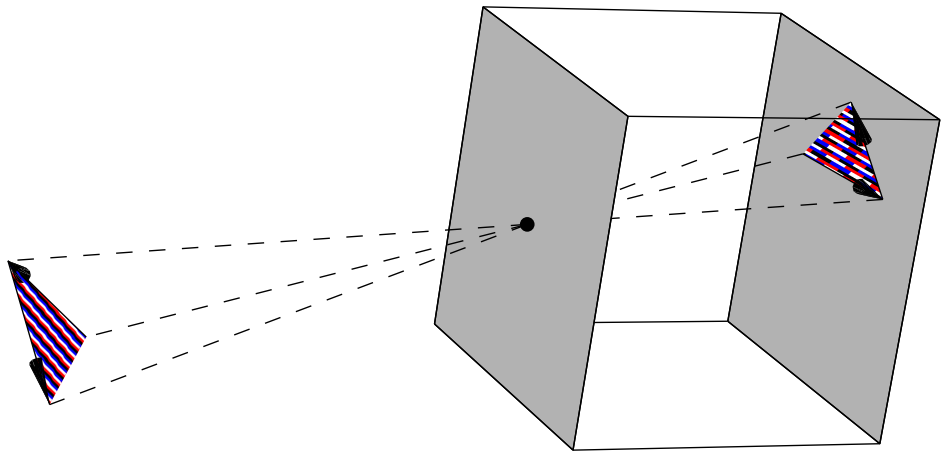

Figure 2.3: The pinhole camera is in principle just a box with a tiny hole in one side. The image is then formed on the opposing wall.

about Leonardo Da Vinci I would be surprised and not so little disappointed if he did not play and experiment with drawing and painting with a camera obscura. And, looking at the central place of the carefully painted mirror-image in the Arnolfini portrait (figure 2.2), vanEycks interest in geometric optics is right in front of our eyes.

The First Pinhole Photographs is attributed to David Brewster, of the Brewster angle fame, in the 1850s. He is also the one that named the pinhole camera pinhole camera. As an imaging device it is ideal in that it has no distorting optical system which also gives almost aberration-free images, instead the resolution is limited by diffraction - due to the small aperture radius of the pinhole. The small aperture also makes the camera very slow requiring long exposures to gather sufficient light - making it more suitable for photography of landscape than sports. The sharpness of an image is determined by the size of the pin-hole and the size to the image surface/plane of the image at the back. Light from a point source will make up a cone through the hole – the smaller the hole the smaller the cone angle and spot size, but only to a limit. The approximate size (for image sharpness) of the optimal pinhole diameter is, $d = 1.9\sqrt{f\lambda}$ When the pin-hole becomes smaller, diffraction comes into play and the ray-optics approximation breaks down. Then the point image will be determined by a combination of the diffraction pattern of the pin-hole and the conical ray bundle. This will be dealt with in section 4.

In order to improve the light gathering properties cameras use lenses to focus light from a larger area. This also reduces the diffraction limited spot size - since the aperture size is increased. However, aberrations will come into action and limit the image resolution.

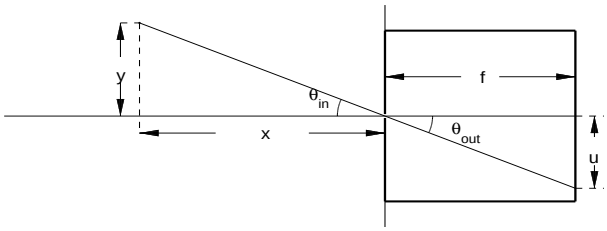

Figure 2.4: The geometry of pinhole imaging.

2.3 The pinhole camera model

Most textbooks on imaging optics use the pinhole camera model and only the pinhole camera model. Its ideal projection characteristics makes it straight forward to apply geometry, and one can show that 4-by-4 matrix algebra perfectly describe the geometry of image formation. This, of course, leads to compact and very neat formalism. However, for real optical systems in general and wider-angle field-of-view optics in particular where there are many lens elements sometimes with large angles of refraction the pinhole model is not guarantied to approximate the imaging geometry well, or even well enough. The exercise in matrix algebra is difficult to extend to the more general case where the optical mapping is not of pinhole type. Therefore we chose a path that is easier to generalise. Occasionally we will return to the pinhole camera model for clean and clear illustrative examples and suitable start points for further excavations.

The imaging geometry of a pinhole camera is particularly simple as can be seen in figure 2.4 outlining the geometry. The image position, u, of an object at (x, y) can be determined by geometry using equilateral triangles:

$$\frac{x}{y} = \frac{u}{f} \Rightarrow u = f \cdot x/y \tag{2.1}$$

For a pinhole camera model this use of equal triangles can be turned into a streamlined and elegant matrix machinery where all of the geometry of the imaging is performed by 4-by-4 matrix algebra. That way all steps, such as translation and rotation of the camera and modification of the focal length, f can be combined into one chain of matrix multiplications. One can equally well determine the horizontal and vertical image position, (u, v) with use of trigonometry. Described that way, light from a direction (ϕ, θ), with θ as the polar angle relative to the optical axis and ϕ as the angle counterclockwise from the horizontal image axis, is projected onto the image according to:

$$(u, v) = f(\tan\theta\cos\phi + u_0, \tan\theta\sin\phi + v_0) \tag{2.2}$$

Where u_0, v_0 is the image coordinate of the optical axis. For a pin-hole camera this does not appear to make much sense, since this formulation does not give as compact a description as the geometry/matrix algebra trail. However, one can modify the simple pinhole camera by centring a half-spherical glass bead just behind

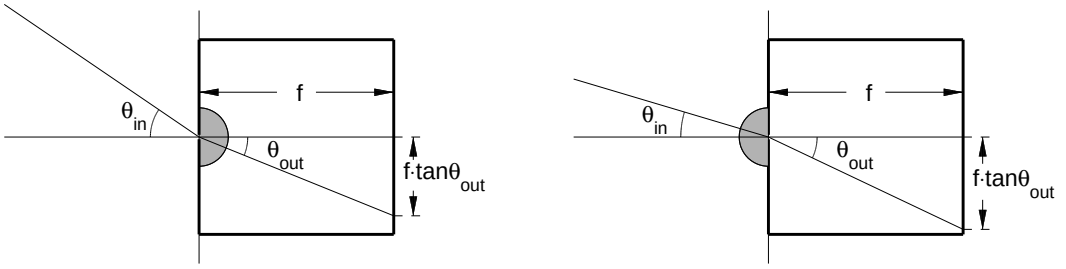

Figure 2.5: The geometry of "beaded" pinhole imaging, with a wide-field-of-view beaded pin-hole camera to the left, where a half-spherical bead, marked in grey, after the aperture refracts rays entering effectively compressing the cone of rays falling on the back plate leading to a wider field-of-view. To the right the bead is positioned in front of the pin-hole aperture, leading to an expansion of the ray-cone by the refraction at the exit, leading to a narrow field-of-view. For both cases the rays passing the pin-hole pass the spherical surface at right angles.

the aperture, as shown to the left in figure 2.5, or just in front as shown to the right. For the wide-field-of-view pin-hole camera to the left the light rays would refract at the flat glass surface at entry but at exit all rays would exit normal to the spherical surface and thus without refraction. The image is still obtained due to the pinhole, but the mapping characteristics is no longer as simple as in equation (2.2). With straightforward use of Snell's law we see that a ray entering the pinhole with an angle θ_{in} is refracted according to

$$\sin\theta_{in} = n_g \sin\theta_{out} \Rightarrow \theta_{out} = \arcsin\frac{1}{n_g}\sin\theta_{in}$$

upon entering the bead. Since the half-spherical bead is centred with respect to the pinhole all rays entering the bead will exit the spherical surface at a right angle and unrefracted, and the projection function is modified to:

$$(u, v) = f\tan(\arcsin\frac{1}{n_g}\sin\theta)(\cos\phi, \sin\phi) + (u_0, v_0) \tag{2.3}$$

For the narrow-field-of-view pin-hole camera, to the right in figure 2.5, the refraction of rays entering through the pin-hole occurs when exiting the bead:

$$n_g \sin\theta_{in} = \sin\theta_{out} \Rightarrow \theta_{out} = \arcsin n_g \sin\theta_{in}.$$

For this case the projection function becomes:

$$(u, v) = f\tan(\arcsin n_g \sin\theta)(\cos\phi, \sin\phi) + (u_0, v_0) \tag{2.4}$$

It is not possible to easily modify the geometry/matrix algebra trail to this more complex imaging characteristics, that is needed to accurately model the imaging characteristics of real lenses, in particular fish-eye lenses and catadioptic systems with wide fields-of-view.

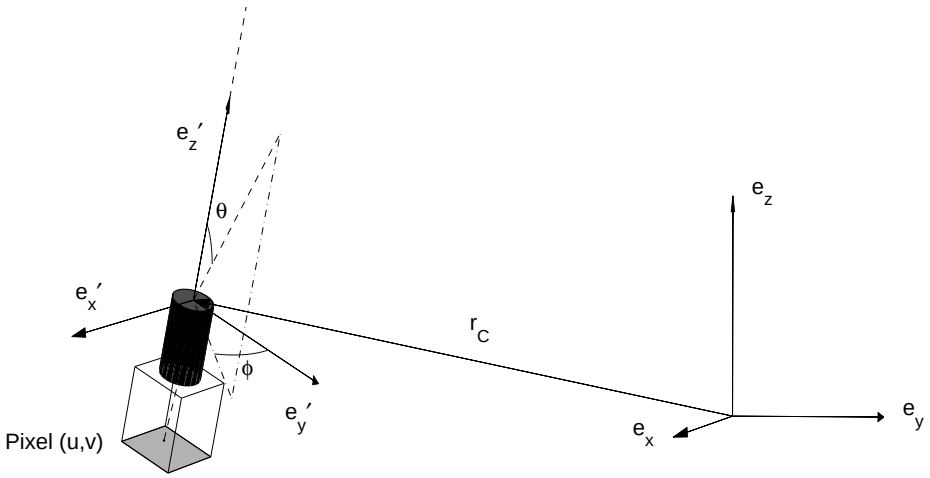

Figure 2.6: The general camera is arbitrarily rotated and has an unknown mapping from line-of-sight direction to image coordinates

2.4 General Camera Models

In general optics used for imaging is built by a number of lenses that are combined to give a desired characteristics - which is often a compromise of optical quality, weight, cost, robustness, size and other parameters. The reason for using lenses in imaging is that it makes us able to collect more light by using a larger aperture over the desired field-of-view. The larger aperture of a lens reduces the size of the Airy-pattern arising from diffraction, that limit the image resolution of pin-hole cameras. But since single lenses with imperfect focus give images with aberration it is better to combine several lenses that combined reduce or even cancel aberrations. Taken together such optical systems can have more than 15 lenses in several groups.

Thus, real cameras in general and cameras with wider field-of-view in particular can have more complicated optical characteristics. Instead of the "geometry trail", we present a path that is easier to generalise. Assume that the camera has a simple projection function:

$$(u, v) = (g(\phi, \theta) + u_0, h(\phi, \theta) + v_0) \tag{2.5}$$

mapping a direction (ϕ, θ) to the image coordinates u, v, as shown in figure 2.6. It is easy to assume something close to rotational symmetry in the optics, which leads to:

$$(u, v) = (f_u f(\theta) \cos\phi + u_0, f_v f(\theta) \sin\phi + v_0) \tag{2.6}$$

Here we have an optical model where θ and ϕ separates with a elliptical symmetry - the scaling factors f_u and f_v allows for non-circularity. The difference in characteristics of course have large impact on pixel line-of-sight, and thus on where and

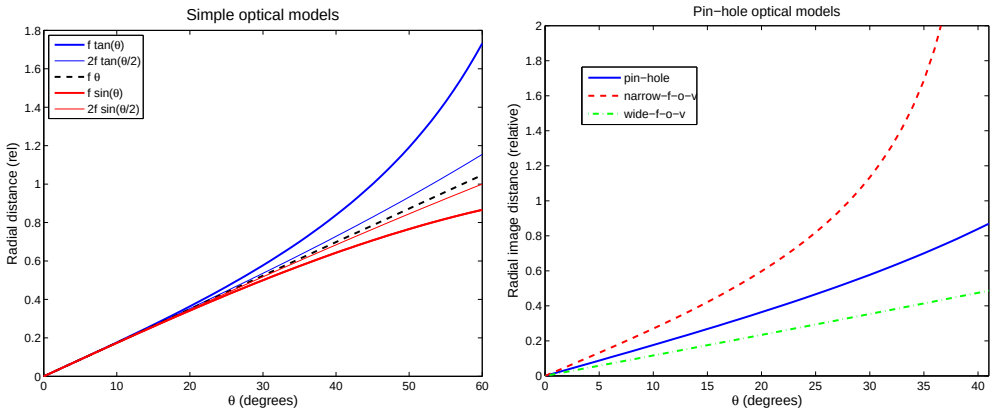

Figure 2.7: Radial distance in image as a function of polar angle for various optical models.

how objects in the camera field-of-view appears in an image but also on pixel field-of-view and thus on optical vignetting. The images of a square grid in figure 2.8 show how the difference between the camera models affect the mapping of real world coordinates. Narrow field-of-view cameras are in general well described by pinhole camera models. Cameras with wider field-of-view can have more complicated optical transfer functions. There are five idealised *radial projection functions* that maps an angle, θ, from the optical axis to a radial distance in the image from the image projection of the optical axis, $((u-u_0)^2+(v-v_0)^2)^{1/2}$,

- Perspective, or pinhole: $f(\theta) = f\tan\theta$
- Stereographic: $f(\theta) = f\tan(\frac{\theta}{2})$
- Sine-law: $f(\theta) = f\sin\theta$
- Equi-solid angle: $f(\theta) = f\sin(\frac{\theta}{2})$
- Equi-distant: $f(\theta) = f\theta$

The difference in geometrical characteristics of these radial projection functions are shown in figure 2.7, where one can see that the difference is mainly seen for larger angles to the optical axis. The perspective and stereographic, and the sine-law and equi-solid angle radial projection functions are pairwise special cases of two more general radial projection functions:

- $f(\theta) = f\tan\alpha\theta$
- $f(\theta) = f\sin\alpha\theta$

Where α now is a shape parameter that controls the projection characteristics of the imaging system. Here it is important to remind oneself that distortion does not degrade the image information, it only changes the mapping of where it appears in the image.

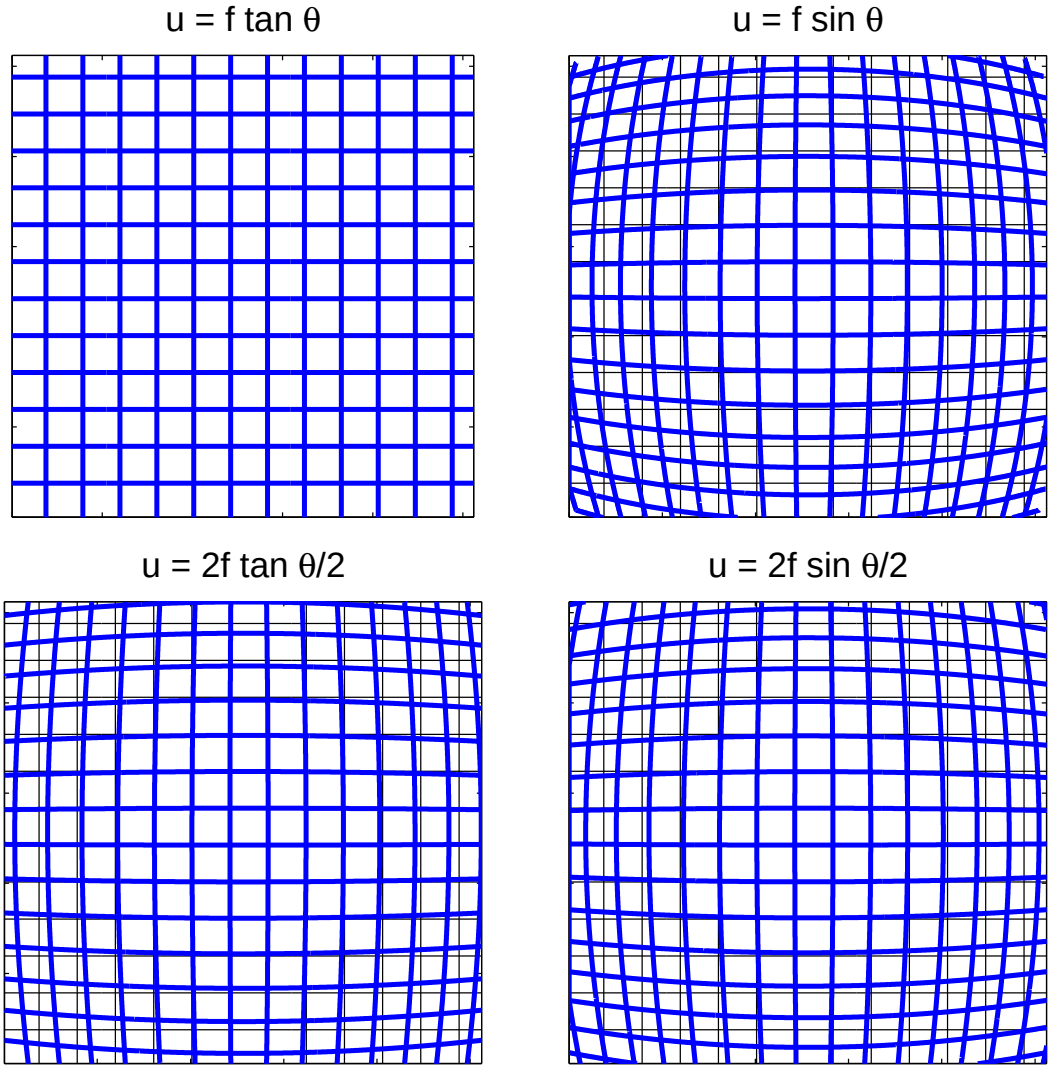

Figure 2.8: Images of a square grid differs depending on the camera projection function, in the top left an image with pinhole projection function, in the top right with a sine-law projection, in the bottom left a stereographic, and in the bottom right an equi-solid angle projection function. The curved projections of lines perpendicular to the optical axis is known as distortion.

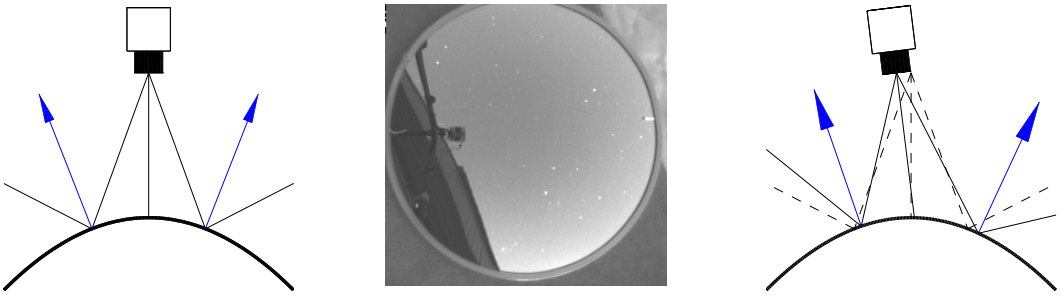

Figure 2.9: For a catadioptic imaging system it is possible to obtain more complex optical transfer functions. With the optical axis of a camera pointing at the central point of a curved mirror, as shown in the leftmost panel, the optical mapping function is symmetric. If there is misalignment, the symmetries are lost and the mapping is more complex.

In principle there is little that constrains a real optical system to any of the models described above and several more complex functions can be used to describe the projection characteristics. Other functions have been used to describe the radial projection functions for camera systems in the field:

- $\alpha\theta + (1-\alpha)\tan\theta$ (ALIS)
- polynomials in θ, $p(\theta)$ (e.g. old ASI)
- rational functions in θ, $r(\theta)/q(\theta)$ (Themis)
- polynomials in $\tan\theta$, $p(\tan\theta)$
- polynomials in $\sin\theta$, $p(\sin\theta)$

Optical imaging systems does not necessarily have to be made with lenses alone, it is possible to combine mirrors and lenses to obtain a desired imaging characteristics. For such more complex imaging systems, for example when the camera is viewing the scene reflected on a curved mirror, as shown in figure 2.9, a so called catadioptic system, there might be no simple closed form solution of the projection function. Provided the mirror has a known analytical shape and is rotationally symmetric; and the camera lens is positioned pointing at the point-of-symmetry it is possible to work out a closed form of the optical mapping function, but as can be seen in figure 2.9, all of these conditions are not necessarily met, sometimes out of design considerations.

The number of free parameters of the camera models depends on the complexity of the radial projection function. But these simple models have four or five "intrinsic" camera parameters:

- u_0, v_0 - image coordinates of the projection of the optical axis.
- f_u, f_v - "focal widths" in the vertical and horizontal direction.
- α - "shape parameter"

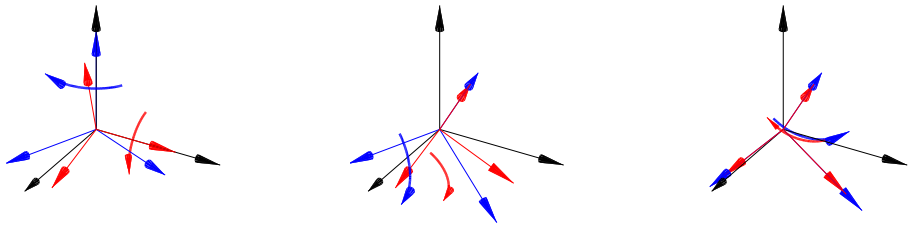

Figure 2.10: Two set of rotations leading to the same final rotation of the coordinate system, the rotations depicted in red denotes "Tait-Bryan angles" and the rotations presented in blue is Euler angles.

More complex camera models can have just any number of free parameters. To determine the pixel field-of-view we have to get these as well as the camera rotations.

Some camera models are easier to work with than others. One preferable property of a camera model is that its inverse is an elementary function or at least simple to calculate. Of the camera models listed above three of the first four have simple inverses, while the $\alpha\theta + (1-\alpha)\tan\theta$ is not.

2.5 Translation and Camera rotations

In the above section we outlined a general method to describe the mapping from a direction (ϕ, θ) to a point (u, v) in the image plane. Here ϕ and θ are the azimuth and polar angle relative to the camera and its optical axis. To be able to calculate the image position of an object, we also need to account for the camera position relative to our "real-world origin" and its rotation.

To uniquely describe the rotation of a rigid body there are a number of different but equally adequate sets of rotations. In figure 2.10 the rotation sequence depicted in blue is described by successive rotations ("Euler angles") around e_z (azimuth), e'_x (zenith) and e''_z (roll):

$$\mathbf{R_1} = \begin{pmatrix} \cos\phi & \sin\phi & 0 \\ -\sin\phi & \cos\phi & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{R_2} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & \sin\theta \\ 0 & -\sin\theta & \cos\theta \end{pmatrix}$$

$$\mathbf{R_3} = \begin{pmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

That give the composite rotation matrix:

$$\mathbb{R} = \mathbf{R_3} \cdot \mathbf{R_2} \cdot \mathbf{R_1} \tag{2.7}$$

But one could equally easy chose to describe the rotations in "Tait-Bryan angles", as if the rotations was performed in a gyro mount, with the first around e''_z, the second rotation around e_y and the third around e'_x :

$$\mathbf{r_1} = \begin{pmatrix} \cos\gamma & \sin\gamma & 0 \\ -\sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{r_2} = \begin{pmatrix} \cos\alpha & 0 & \sin\alpha \\ 0 & 1 & 0 \\ -\sin\alpha & 0 & \cos\alpha \end{pmatrix}$$

$$\mathbf{r_3} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\beta & \sin\beta \\ 0 & -\sin\beta & \cos\beta \end{pmatrix}$$

That give the same composite rotation matrix:

$$\mathbb{R} = \mathbf{r_3} \cdot \mathbf{r_2} \cdot \mathbf{r_1} \tag{2.8}$$

In practice it is advisable to use a set of rotation angles that reflect the design of the hardware. The choice in ALIS to mount the cameras in a gyro-mount makes the natural choice for describing the rotations the "Tait-Bryan angles" to describe the rotations. But in principal one could without loss of accuracy have used Euler angles. As shown in figure 2.10, the final rotation of the camera is identical, while the rotation angles and the order of the rotations are not the same. Case in point: In the ALIS project there seems to have been a mistake in the order of the first two rotations somewheres between the master-plan, controller software, camera positioning system and analysis software. This is no problem for the data analysis as long as we use angles estimated by calibration from images — since this calibration use the same set of rotations as everywhere in the analysis.

Translations are pretty trivial to handle:

$$\bar{R}' = \bar{R} - \bar{r}_0 \tag{2.9}$$

2.6 Calibration of camera parameters

For studies of the atmosphere and the ionosphere it is appealing to use calibration images made under the same conditions with exactly the same rotations and use background stars as a way of determining the lines-of-sight. In this way the refraction of light in the atmosphere is automatically accounted for.

The requirement that the determination of the viewing directions should be accurate to within 0.02° implies that the sky positions of these stars need to be known to within 0.02°. This accuracy requirement demands that the sky positions of the stars have to be rigorously corrected for precession when calculated from the star catalogue. The star catalogue used in this work is the "Bright Star catalogue" (BSC,

formerly known as the "Yale Bright Star" catalogue, YBS) which contains some 9000 selected stars down to magnitude 6.8. In order to verify the accuracy of the position calculations, the results have been compared with the star chart program "xephem" (available at the time of writing at http://www.clearskyinstitute.com/xephem/). It is not necessary to correct for the nutation, the minor variation in the inclination of the earth axis due to the torque from the moon. Likewise it is not necessary to include the corrections for aberration, proper motion and parallax in the calculations. These higher order corrections are smaller than 0.01°.

To proceed with determining the pixel lines-of-sight it is necessary to be able to identify a large number of stars in an image, preferably more than 100 evenly distributed over the image plate. For these "star pixels" we know the lines-of-sight to within the required accuracy. For the intermediate pixels we can choose either to interpolate between the lines-of-sight from neighbouring star pixels or use an optical transfer function that describes where the light from a direction in space $\bar{e}(\phi, \theta)$ falls on the image. The identified stars are used to adjust the parameters of the optical transfer function so that the calculated directions of the stars are mapped down to their respective points on the image.

For the auroral and atmospheric studies which are our main interest, where the aim is to determine the position of objects at altitudes between 15 and 500 km, the required parameter is the path of the light through the atmosphere.

Due to the variations in density and temperature with altitude, the refractive index of the atmosphere varies from 1 in free space to approximately 1.0003 at sea level. This variation in refractive index causes an object outside the atmosphere to appear closer to zenith than it actually is. An everyday example of this is that when the sun disappears from the horizon its centre is already 0.55° below the horizon.

For a flat plane parallel atmosphere the total change in zenith angle is easily calculated by Snell's law:

$$n_0 \sin(\theta_0) = n_1 \sin(\theta_1) = \cdots = n_\infty \sin(\theta_\infty) \tag{2.10}$$

where n_i is the refractive index and θ_i is the zenith angle at altitude i. From this it is straightforward to calculate the difference between the true zenith angle, θ_∞, and the apparent zenith angle, θ_0, where

$$\theta_0 = \arcsin(\frac{n_\infty}{n_0} \sin \theta_\infty) \tag{2.11}$$

This simple relation is correct to within 0.001° for zenith angles up to 60°; for larger zenith angles it is not possible to neglect the curvature of the earth. A more general correction that is good for zenith angles up to 75° is [*Smart*, 1977]:

$$\theta_0 = \theta - 58''.16 \tan \theta + 0''.067 \tan^3 \theta \tag{2.12}$$

In order to determine the path through the atmosphere there are three choices: to use the apparent zenith angle and to perform a path tracing in order to calculate the true path of the light, to use the apparent zenith angle and assume that the

light travels in a straight line, or to use the true zenith angle and use that angle for calculating the line-of-sight.

If we use the true zenith as the zenith angle of the light ray there will be an error between that ray and the true path of light due to the refraction in the atmosphere as can be seen in Figure 2.11, for which the ray paths are plotted from the ground to

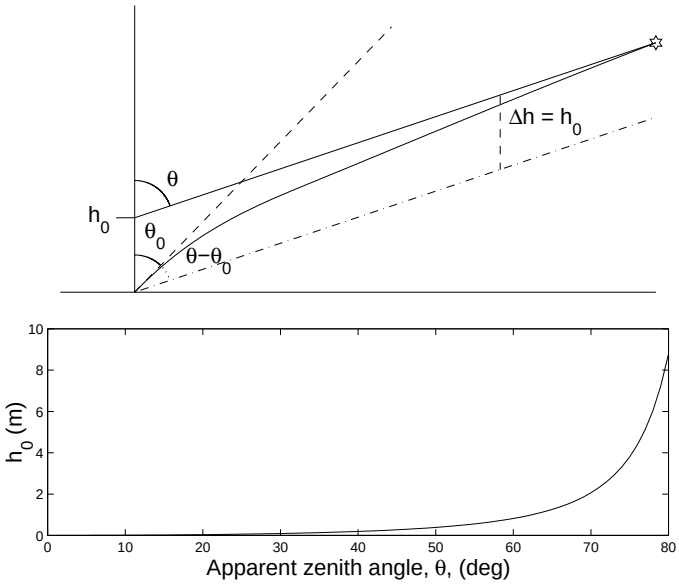

Figure 2.11: Ray paths through the atmosphere and the correction in altitude between true zenith ray path and true zenith line-of-sight.

the sky. The error will be equivalent to an increase in the altitude of the observing site. This displacement error in the altitude of the observation site, h_0, can be expressed as

$$h_0 = R_0 \cdot (\frac{n_0 \cdot \sin z_0}{\sin(z_0 + (n_0 - 1 \cdot \tan z_0))} - 1) \tag{2.13}$$

For zenith angles of up to 70° this accounts for a negligible error of less than ≈ 1.5 m, [*Green*, 1985].

If we were to use the apparent zenith angle, the error between the inferred and the true paths of light would be:

$$\Delta r = |\bar{r}| \cdot (\theta_0 - \theta_\infty) \tag{2.14}$$

For a zenith angle of 45° the difference between true and apparent zeniths is approximately 1 arc minute which gives an error of approximately 60 m at 200 km distance. Here it can be concluded that use of the true zenith angles and approximations of the light path with a straight line gives errors that are within the acceptable accuracy. The small error that results from this approximation can be compared with all other errors in the analysis and weighted against the computationally heavy and cumbersome task of performing a path-tracing through an atmosphere whose refractive index structure as a function of altitude is, at best, only roughly known.

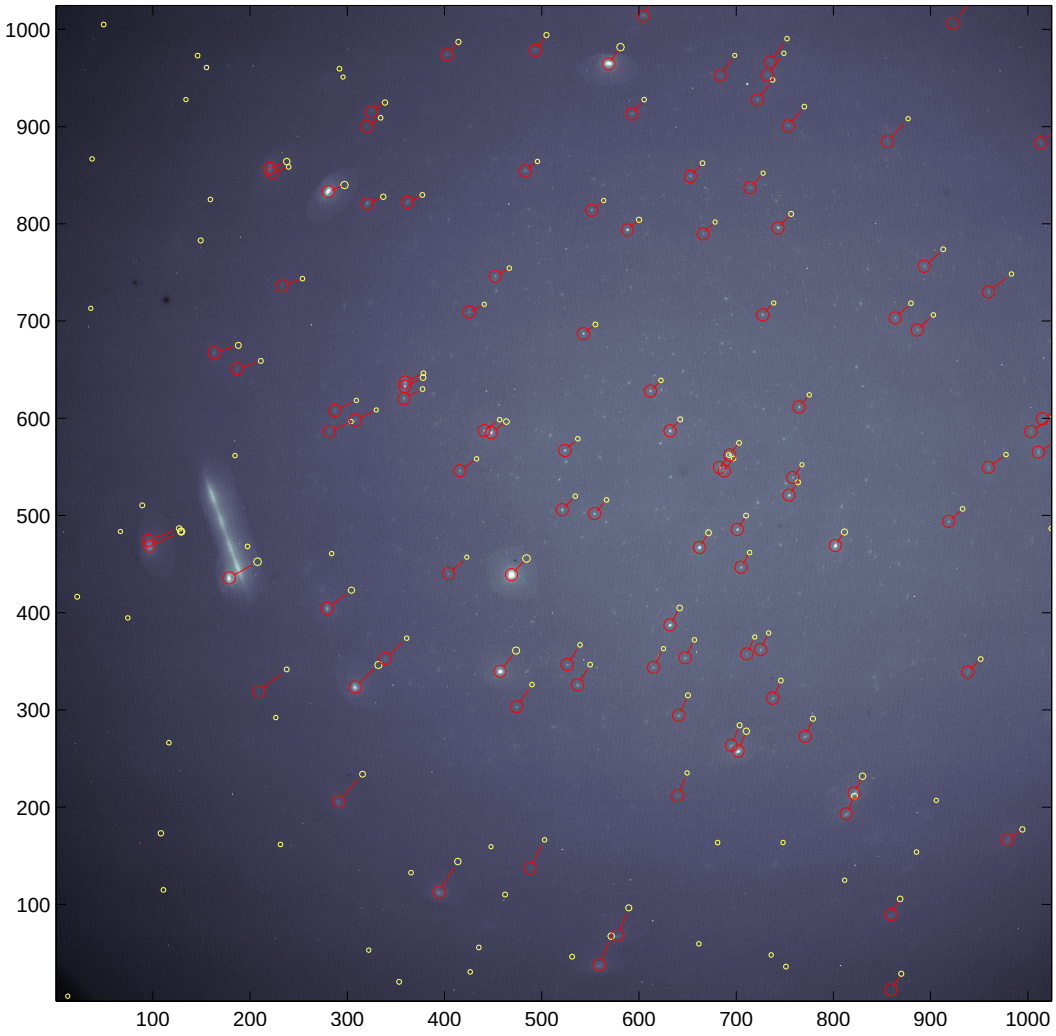

Figure 2.12: Image with overlaid star-field, shown in yellow, calculated for non-optimal camera parameters for selected stars shown in red.

One way to obtain the geometric calibration is to image a set of points in well known directions (ϕ_i, θ_i). When these points are identified in an image we have observed the image coordinates (u_i, v_i); and the camera model will give $(u_i, v_i) = (f_u(\phi_i, \theta_i, \bar{p}), f_v(\phi_i, \theta_i, \bar{p}))$ where $\bar{p}$ is the camera parameters to be determined. Then the camera parameters are adjusted so that $(u_i, v_i) - (f_u f(\theta_i) \cos\phi_i + u_0, f_v f(\theta_i) \sin\phi_i + v_0)$ has no systematic patterns and the sum of the squared residuals

$$\sum_i (u_i - (f_u f(\theta_i) \cos\phi_i - u_0)^2 + (v_i - f_v f(\theta_i) \sin\phi_i - v_0)^2 \qquad (2.15)$$

is minimised. A remaining question is which camera model to use, and how many parameters to fit for.

The answer is: the simples camera model that gives a good fit, and as many parameters as needed to remove all systematic residuals.

The ultimate limit is to get away with all systematic errors between the projected positions of our calibration targets (stars) and their actual image position. And

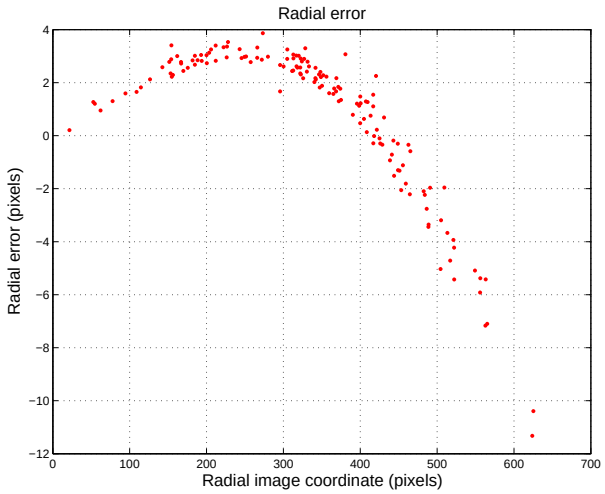

Figure 2.13: This is the variation of errors in the radial direction as a function of radial position for an ALIS camera, assuming pinhole optics, $f \tan\theta$. There are as we can see obvious systematic errors.

get the error scatter to within a suitably small fraction of a pixel.

With a catadioptic system, that for the sake of argument is misaligned, as in the rightmost panel in figure 2.9, there will be a large number of parameters to determine, the intrinsic camera parameters $(f_u, f_v, u_0, v_0, \alpha, \cdots)$, the position (r_C) and rotation $(\phi_{C1}, \phi_{C2}, \phi_{C3})$ of the camera relative to the mirror, and then the parameters corresponding to the intrinsic parameters of the mirror (that determine its curvature, the number depending on its shape, for simplicity, ρ_u, ρ_v) and its rotation $(\phi_{M1}, \phi_{M2}, \phi_{M3})$. This increase in the number of parameters to determine makes accurate calibration of a parametrised camera model more cumbersome, first in making an explicit model of the geometries involved and then in estimating the parameters. There is an alternative approach, that circumvents these difficulties. With a large number of calibration points in known directions, the stars in the examples here, one can also opt for direct interpolation of the geometric projection. Then the objective is to find suitable fitting surfaces mapping between the line-of-sight directions and image coordinates. Surface fits of the azimuth, ϕ, and polar angles, θ, to the optical axis have two pitfalls: firstly the function for azimuthal angles has to be 2π-periodic, and secondly the the polar angle from the optical axis will be approximately proportional to $|\theta|$ for small θ. To make the azimuthal surface 2π-periodic it is necessary to use the identified calibration points with azimuthal angles close to but smaller than π close to the azimuthal cut at $\pm\pi$ also at angles smaller than $-\pi$ (or at 0–2π, whichever we choose). The cusp-like character of $|\theta|$

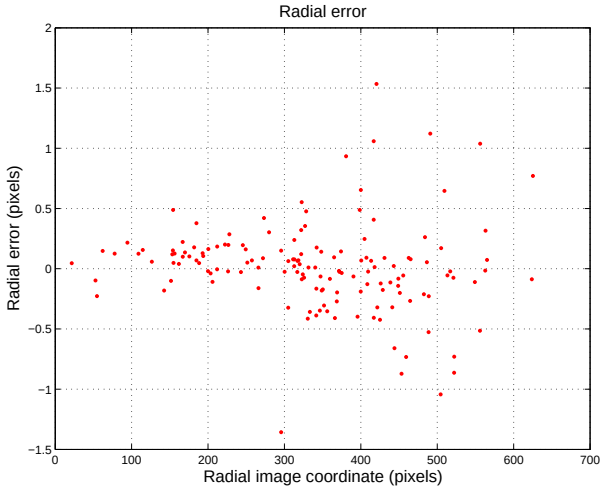

Figure 2.14: This is the same plot but now for the slightly more complex optical transfer function: $(1-\alpha)\tan\theta + \alpha\theta$. Here there are no obvious systematic errors.

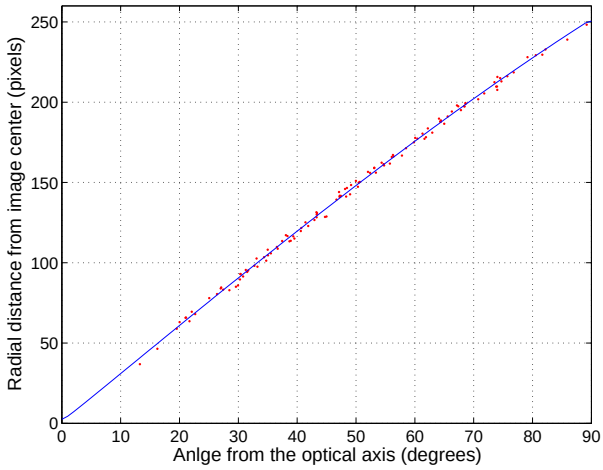

Figure 2.15: For a sufficiently optimal radial transfer function there is no obvious systematic of the residual between observed and calculated radial position of the calibration points, as shown in the example above for the NIPR ASC.

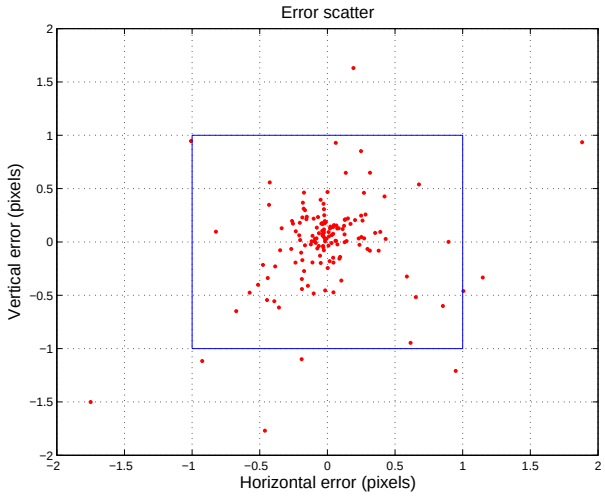

Figure 2.16: Two-dimensional error scatter showing the fit between the stars as found in the image and projected from the stellar position in the sky through the "model-optics". The fit is good since most stars are within ±1 pixel.

makes it difficult to get a proper fitting surface since it will be an extrapolation from the smallest θ_{cal} to $\theta = 0$ - since the fitting surface also has to be a smoothing surface not to be effected by noise. A more tractable set of fitting surfaces describing the mapping between the image coordinates and the line-of-sight directions is between the components of the line-of-sight unit vectors and the image coordinates. The variation of $\bar{e}_{pix} = [e_1, e_2, e_3]$ will all be smooth and well-behaved as functions of image coordinate: $e_1(u, v), e_2(u, v), e_3(u, v)$. Thus it is straightforward to calculate the line-of-sight unit vectors for all identified calibration points, make a fitting surface from image coordinates (u, v) to each component, and inverse fitting surfaces from line-of-sight components to image coordinates ($u(e_1, e_2, e_3)$ and $v(e_1, e_2, e_3)$), as shown in figure.

2.7 Accuracy estimates

In order to know how far we can reach in retrieving spatial information from our images after the calibration we need to make some accuracy estimates. It is of some importance to know how our accuracy (or rather lack of...) in pixel line-of-sight propagates through further analysis. In order to estimate this we can calculate the error in the triangulated position of an imaged object as a function of line-of-sight uncertainties:

Triangulation error

Here, a first order error analysis is easily accomplished if we look at stereoscopic triangulation of an object point found in two images from sites at $\bar{r}_1$ and $\bar{r}_2$.

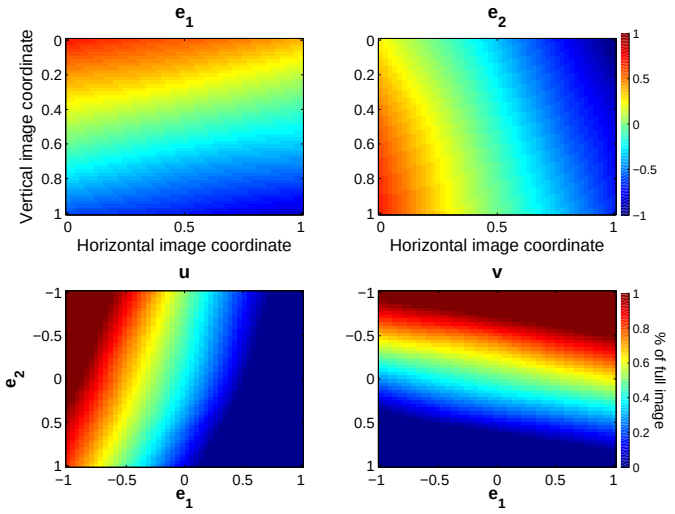

Figure 2.17: Interpolants for the catadioptic setup displayed in figure 2.9, in the top left panel is the interpolating surface for horizontal image coordinates as a function of the line-of-sight vector components in West-East ($\mathbf{e_1}$) and South-North ($\mathbf{e_2}$) directions. In the top right panel is the interpolating surface for the vertical image coordinate. In the bottom panels are the interpolating surfaces for the inverse mapping - from image coordinates (u, v) to the horizontal line-of-sight unit vector components.

If we let the two lines-of-sight Γ_1, Γ_2 be described by their angles Φ_1, Φ_2 relative to the vector $\bar{r}_1 - \bar{r}_2$, as shown in figure 2.18, it is straightforward to calculate the uncertainty in the position of $\bar{r}_{point}$ due to errors in the line-of-sight parameters $\Delta\Phi_1, \Delta\Phi_2$ and the station separation Δl.

Simplifying the problem to two dimensions we get

$$h = l\frac{\tan\phi_1 \cdot \tan\phi_2}{\tan\phi_1 + \tan\phi_2}$$
$$x = l\frac{\tan\phi_1}{\tan\phi_1 + \tan\phi_2}$$

for the position of the top corner. The sensitivity of h and x to errors in l, ϕ_1, and ϕ_2 is then estimated by the full differentials:

$$dh = \frac{\partial h}{\partial l}\Delta l + \frac{\partial h}{\partial \phi_1}\Delta\phi_1 + \frac{\partial h}{\partial \phi_2}\Delta\phi_2$$
$$dx = \frac{\partial x}{\partial l}\Delta l + \frac{\partial x}{\partial \phi_1}\Delta\phi_1 + \frac{\partial x}{\partial \phi_2}\Delta\phi_2$$

which can be expanded by calculating the partial derivatives (the ambitious student can check my alge-

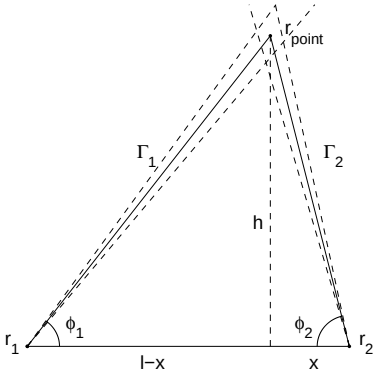

Figure 2.18: Triangulation error triangle showing the error propagation from uncertainties in the line-of-sight directions to uncertainty of the triangulated position.

	Case 1	Case 2	
$\partial h/\partial l$	0.3889	1.3268	()
$\partial h/\partial \phi_1$	0.5280	2.0484	$(km/^\circ)$
$\partial h/\partial \phi_2$	0.2249	1.5841	$(km/^\circ)$
$\partial x/\partial l$	0.3264	0.2340	()
$\partial x/\partial \phi_1$	0.4431	0.3612	$(km/^\circ)$
$\partial x/\partial \phi_2$	-0.3896	-0.9146	$(km/^\circ)$

Table 2.1: Accuracies required to achieve designated triangulation errors.

bra):

$$\frac{\partial h}{\partial \phi_1} = \frac{l \tan^2 \phi_2 \left(1 + \tan^2 \phi_1\right)}{\left(\tan \phi_1 + \tan \phi_2\right)^2} \quad \frac{\partial x}{\partial \phi_1} = \frac{l \left(1 + \tan^2 \phi_1\right) \tan \phi_2}{\left(\tan \phi_1 + \tan \phi_2\right)^2}$$

$$\frac{\partial h}{\partial \phi_2} = \frac{l \tan^2 \phi_1 \left(1 + \tan^2 \phi_2\right)}{\left(\tan \phi_1 + \tan \phi_2\right)^2} \quad \frac{\partial x}{\partial \phi_2} = -\frac{l \tan \phi_1 \left(1 + \tan^2 \phi_2\right)}{\left(\tan \phi_1 + \tan \phi_2\right)^2}$$

$$\frac{\partial h}{\partial l} = \frac{\tan \phi_1 \tan \phi_2}{\tan \phi_1 + \tan \phi_2} \qquad \frac{\partial x}{\partial l} = \frac{\tan \phi_1}{\tan \phi_1 + \tan \phi_2}$$

From these equations it is possible to calculate the propagation of uncertainties in the lines-of-sight to the positions for all directions.

As a rule of thumb, as the station separation is small compared to the distance to the object, i.e. $|\bar{r}_1 - \bar{r}_2| \to 0$ the sensitivity in the determination of positions to errors in the direction of the lines-of-sight increase rapidly.

(Humans use stereoscopy only within a few meters. Further away spatial information is inferred from perspective object size and experience. Also same reason distance to stars cannot be determined by stereoscopy except for the closest.)

In table 2.1 are two examples with typical values for stratospheric (case 1) and auroral (case 2) work.

To get the estimated distances correct to within 50 m, we need to know the base line, Δl, within ± 50 m and the line-of-sight directions, $(\Delta\phi_1, \Delta\phi_2)$, to within $\pm 0.02^\circ$. For stratospheric studies, where the parallax is larger, the requirements are less strict.

The analytically inclined can readily extend this analysis to the non-planar case.

An easier way to estimate the accuracy of triangulation can be made with a straight-forward Monte Carlo simulation. Then one can instantly calculate the distribution of the residual between the true point and the triangulated point. In FIGURE 2.19 a set of such calculations show the spread of triangulated positions from the "true" position for the case when the standard deviation of the line-of-sight direction for selected pixels are 0.03°, with the "cameras" located 100 km apart. Here one should note that this uncertainty also includes the accuracy of the manual selection of pixels.

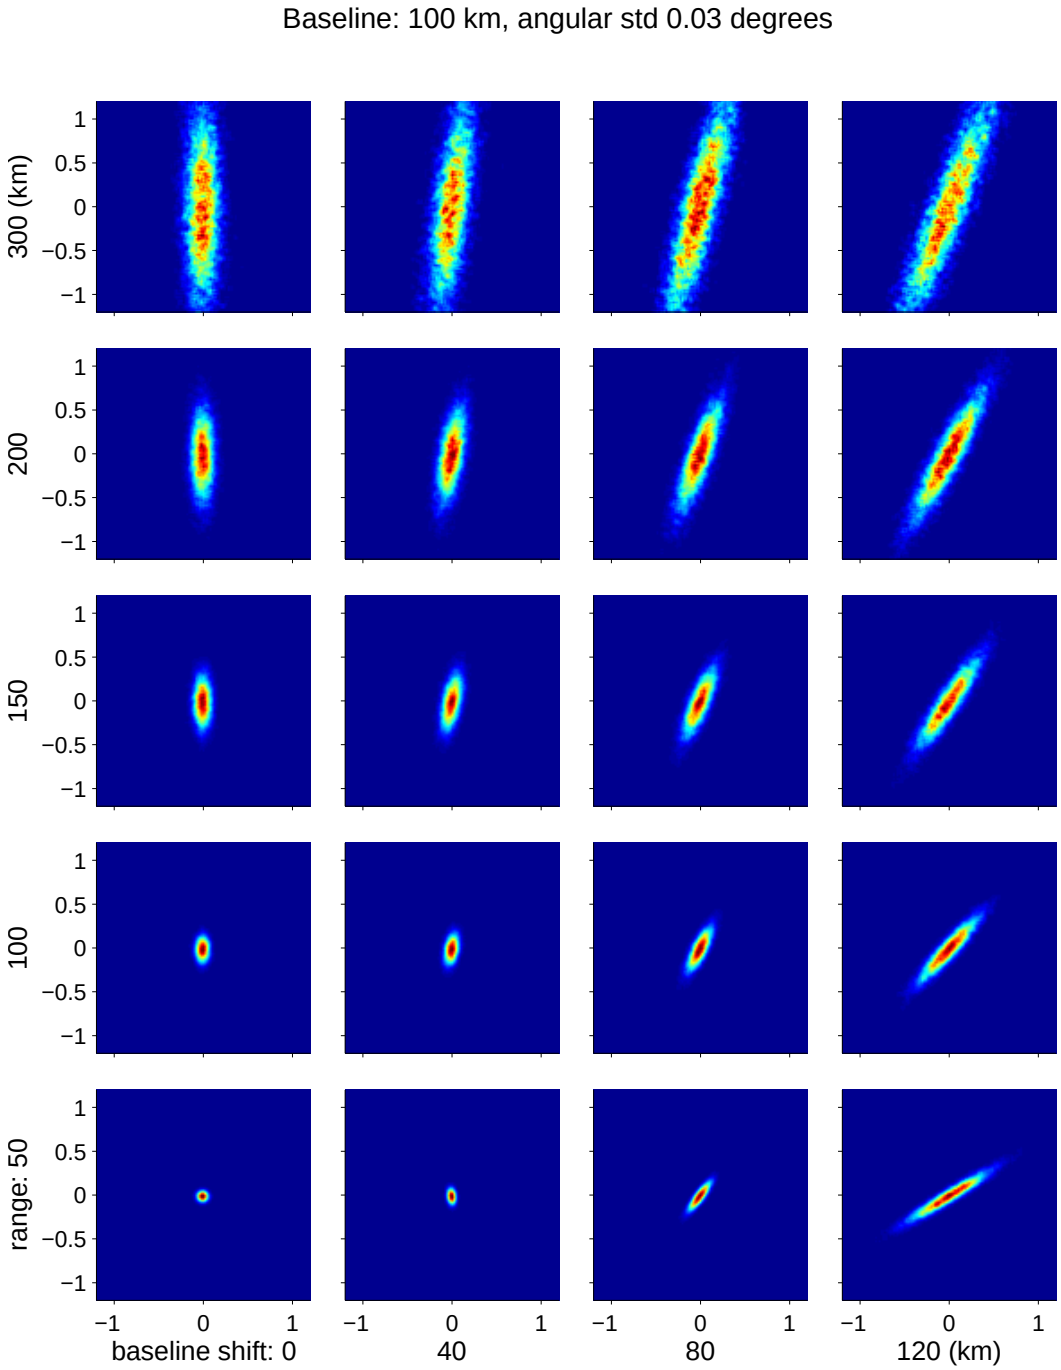

Figure 2.19: Results of Monte Carlo calculations of stereoscopic triangulation accuracy show that the accuracy are easily within 0.2 km in the direction parallel to the baseline, and within ±1 km in range. This holds for this set of points located between 50 and 300 km away and 0 to 120 km shifted from the central point along the baseline between the imaging sites.

2.8 Summary

In this chapter we have outlined the geometry of the imaging process. Modelling the geometrical aspects of imaging can be summarised in the following steps:

- Translate the real-world coordinates to the un-rotated camera coordinate system ($\bar{R}' = \bar{R} - \bar{r}_{camera}$)
- Rotate the translated real-world coordinates to the rotated system of the camera ($\bar{R}'' = \mathbb{R} \cdot \bar{R}'$).
- Convert these rotated-translated coordinates to spherical coordinates $(X'', Y'', Z'') \rightarrow (\Phi, \Theta, \rho)$.
- Use the projection function that describes your camera to calculate the image positions of the objects inside the camera field-of-view ($u = (f_u f(\theta) \cos\phi + u_0, v = f_v f(\theta) \sin\phi + v_0)$)

These are the fundamental steps of image formation. We need to understand them with sufficient depth to be able to model the geometry of the imaging process.

Chapter 3

Irradiance response and flat-field correction

...or why we get nonuniform illumination

> *There can be no transforming of darkness into light and of apathy into movement without emotion. – Carl Gustav Jung*
>
> *I will love the light for it shows me the way, yet I will endure the darkness for it shows me the stars. Og Mandino*
>
> *Darkness cannot run out darkness, only light can do that; Hate cannot run out hate, only love can do that. – Dr. Martin Luther King Jr.*
>
> *Education is the movement from darkness to light. – Allan Bloom*

3.1 Introduction

In the previous chapter we outlined the geometric properties of the imaging process. This makes it possible for us to calculate the image position of objects, but not their brightness in the image. In this chapter we will take this next step in image formation. We will outline how to calculate the variation in image intensity.

Even if we have an isotrope and homogeneous flux of photons in front of the lens we will not get a flat image with constant intensity in all pixels. The darkening towards the image corners is known as vignetting, and there are two sources contributing. One source is mechanical vignetting, caused by reduction in limiting aperture of the optics at increasing angles relative to the optical axis. Natural vignetting is caused by a combination of the variation in pixel field-of-view and effective collecting area of the optics. Natural vignetting for optical systems with pin-hole characteristics are approximately proportional to $\cos^4\theta$, which often is stated as a "law of nature" for vignetting. This is not really true. Here we will derive a general (but slightly idealised) description of the image intensity variation from a uniform and isotropic source. We will also outline a method to correct for local variation of sensitivity on the detector, or photo-response non-uniformity.

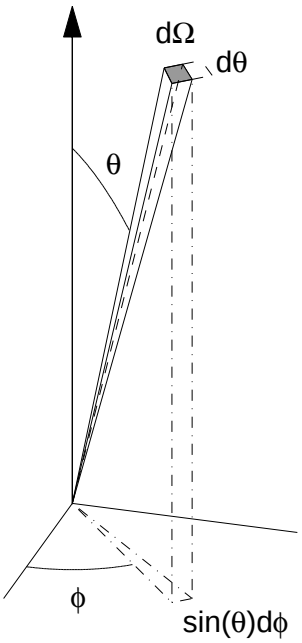

Figure 3.1: Definition of the pixel field-of-view solid angle and related quantity.

3.2 Natural vignetting

Here we will derive a general (but slightly idealised) description of the variation in the number of photons striking the detector surface. The cause of this variation in image intensity is from two main factors, the variation in solid angle of the pixel fields-of-view and the apparent area of the front lens in the direction of the fields-of-view of the pixels.

Given an optical transfer function as outlined in section 2.4 above, the pixel field-of-view, $d\Omega$, can be calculated by standard calculus, once the parameters of the optical transfer function are determined. With θ as the angle relative to the optical axis of the camera, and ϕ as the azimuthal angle, as plotted in figure 3.1, the field-of-view solid angle, $d\Omega$, is

$$d\Omega(u,v) = \sin\theta d\phi d\theta \tag{3.1}$$

where it is important to remember that ϕ and θ are functions of pixel coordinates (u,v) in the image – effectively the inverse of the optical mapping function, equation (2.5). To calculate the solid angle for the field-of-view of one pixel (u,v) we have to calculate $d\phi d\theta$, which can be calculated in the same way as variable transformations are performed in calculus:

$$\sin\theta d\phi d\theta = \sin\theta |\frac{\partial(\phi,\theta)}{\partial(u,v)}| du dv \tag{3.2}$$

Since the inverse of the optical mapping function might be tricky to obtain analyti-

cally it is more convenient to express $d\Omega$ as

$$d\Omega(u,v) = \sin\theta |\frac{\partial(u,v)}{\partial(\phi,\theta)}|^{-1} dudv \tag{3.3}$$

for which (u,v) are the horizontal and vertical image coordinates and $|\frac{\partial(u,v)}{\partial(\phi,\theta)}|$ is the absolute value of the Jacobian determinant of the optical transfer function.

Further we have to take the variation in effective collecting area of the optics into account. In the simplest approximation, the area of the front lens as seen from a direction varies with the polar angle θ relative to the optical axis as:

$$A(\theta) = A_L \cos\theta \tag{3.4}$$

Here A_L is the area of the front lens or the limiting aperture of the optics at normal incidence and θ is the angle relative to the optical axis.

Combining equations (3.3 and 3.4) we get a general expression (ignoring transparency variations in optics and mechanical vignetting for now) for the image intensity variation from uniform isotrope radiance:

$$I_{flat} \propto A_L \cos\theta \sin\theta \left|\frac{\partial(u,v)}{\partial(\phi,\theta)}\right|^{-1} dudv \tag{3.5}$$

When using this expression for corrections it is important to remember that multiplication of image intensity might change the relation between signal and noise.

For optics with rotational symmetry (or with eigenvectors $\parallel$ to the image coordinates) where $(u,v) = (f_u f(\theta)\cos\phi + u_0, f_v f(\theta)\sin\phi + v_0)$, we can simplify further to obtain

$$I_{flat} \propto \frac{\cos\theta \sin\theta}{|f_u f_v f(\theta) \frac{\partial f(\theta)}{\partial\theta}|} \tag{3.6}$$

To summarise, figure 3.2 shows a number of flat-field curves for some of the optical transfer functions listed before.

3.3 Examples of flat-field correction calculation

For pin hole type optics where $(u,v) = f(\tan\theta\cos\phi, \tan\theta\sin\phi)$ the flat field variation is

$$I_{flat} = \cos\theta\sin\theta \left|\frac{\partial(u,v)}{\partial(\phi,\theta)}\right|^{-1} = \cos\theta\sin\theta \left|\frac{\partial(f(\tan\theta\cos\phi, \tan\theta\sin\phi))}{\partial(\phi,\theta)}\right|^{-1} \tag{3.7}$$

Expanding the Jacobian determinant leads to

$$\frac{\partial(f(\tan\theta\cos\phi, \tan\theta\sin\phi))}{\partial(\phi,\theta)} = \begin{vmatrix} -f\tan\theta\sin\phi & f\cos^{-2}\theta\cos\phi \\ f\tan\theta\cos\phi & f\cos^{-2}\theta\sin\phi \end{vmatrix} = f^2\sin\theta/\cos^{-3}\theta$$

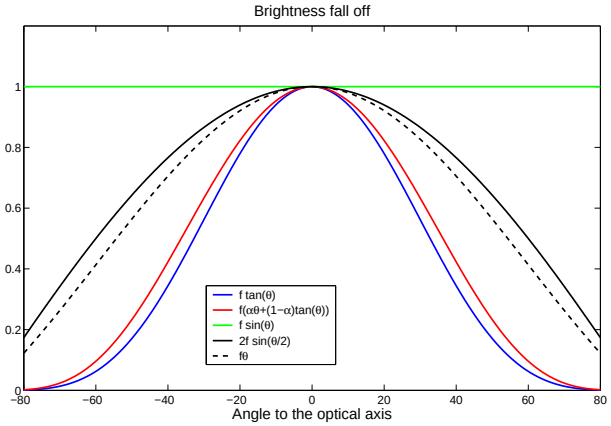

Figure 3.2: Flat field response curves for 5 different optical transfer functions.

Inserting this into equation (3.7) gives

$$I_{flat} = \frac{1}{f^2}\frac{\cos\theta\sin\theta\cos^3\theta}{\sin\theta} = \frac{1}{f^2}\cos^4\theta \tag{3.8}$$

which as expected agrees with the $\cos^4$-law presented in most textbooks. This is, *again*: not a generally true result. If we for example have a fish-eye lens with an optical mapping function $(u, v) = f(\sin\theta/2\cos\phi, \sin\theta/2\cos\phi)$ we get

$$I_{flat} = \cos\theta\sin\theta\left|\frac{\partial(u,v)}{\partial(\phi,\theta)}\right|^{-1}dudv = \cos\theta\sin\theta\left|\frac{\partial(f(\sin\frac{\theta}{2}\cos\phi, \sin\frac{\theta}{2}\sin\phi))}{\partial(\phi,\theta)}\right|^{-1} \tag{3.9}$$

where the Jacobian determinant expands to

$$\frac{\partial(f(\sin\frac{\theta}{2}\cos\phi, \sin\frac{\theta}{2}\sin\phi))}{\partial(\phi,\theta)} = \begin{vmatrix} -f\sin\frac{\theta}{2}\sin\phi & f/2\cos\frac{\theta}{2}\cos\phi \\ f\sin\frac{\theta}{2}\cos\phi & f/2\cos\frac{\theta}{2}\sin\phi \end{vmatrix} = \\ = \frac{f^2}{2}\sin\frac{\theta}{2}\cos\frac{\theta}{2} = \frac{f^2}{4}\sin\theta$$

When this is inserted in equation (3.9) this now leads to

$$I_{flat} = \frac{4}{f^2}\frac{\cos\theta\sin\theta}{\sin\theta} = \frac{4}{f^2}\cos\theta \tag{3.10}$$

The example shown in figure 3.3 illustrates the application of the correction outlined above on ALIS data, taken with a lens system with 55° field-of-view. ...which apparently is good enough for government work.

So far we have ignored mechanical vignetting (that is variations in clear lens aperture A_L with angle θ to the optical axis) and variation of transmission through the optics. But when we compare the result of this flat-field correction for 55° and 75° field-of-view images from ALIS we see (in figure 3.4) that towards the corners

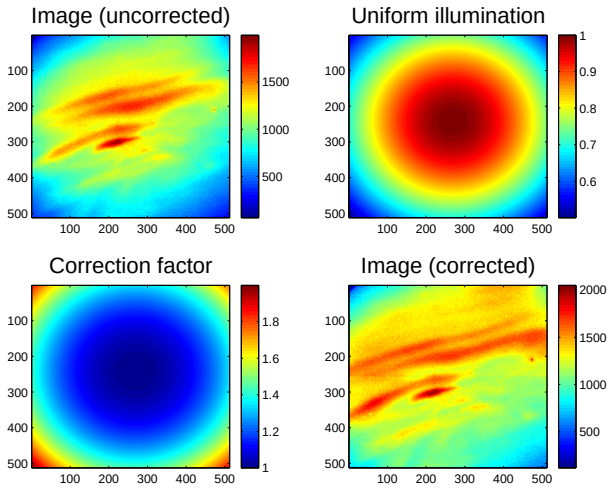

Figure 3.3: Flat-field correction procedure of ALIS data. Top left panel show the uncorrected ALIS image, the top right panel displays the image calculated from a flat radiance. In the bottom left the correction factor, which is the inverse of the top right is shown, and the final flat-field corrected image is displayed in the bottom right panel

of the 75° image and to some extent at the edges this correction is not perfect. There might be several factors contributing to this. One is increased losses from reflection in the lens surfaces at higher angles of incidence. The other effect, that comes into play for more complex optical systems with a large number of lenses and apertures, is the effects of mechanical vignetting – the decrease of effective aperture with increasing angle relative to the optical axis. The higher order corrections that are needed could in principal be determined through accurate ray tracing but must in practice be determined from measurements.

3.4 Mechanical vignetting

For a single lens camera all ray-bundles should have the effective exit area proportional to the entrance area - as described by equation (3.4). For larger optical systems with numerous lenses and apertures this might not be the case. If we take a closer look at a pin-hole where the aperture diameter is of the same order as the thickness of the front plate we see, in figure 3.5, that only the ray-bundle parallel to the optical axes will be unobstructed The relative reduction in unobstructed area depend on the relation between the radii of the apertures and the distance between them; it is even possible to obtain the relation in closed form - but that is a little too long to type and read. So in figure 3.6 we only present the relation for some typical numbers. For a real optical system it is more complicated since ray-bundles also undergo refraction between the limiting apertures - and there might also be more than two effective apertures. Taking all this into account gets increasingly tedious and cumbersome. Meticulous and careful calibrations might be a more successful way to correct for severe mechanical vignetting, however, when the reduction is large

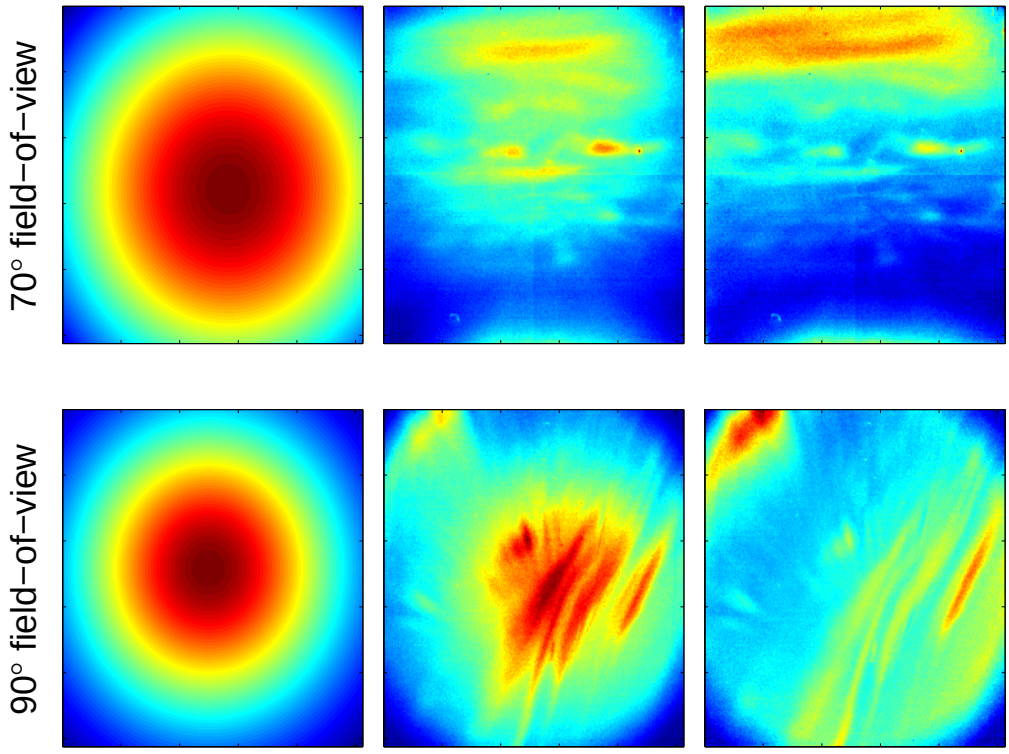

Figure 3.4: Left column shows the flat-field correction images for ALIS camera # 1 (top row) and # 2 (bottom). In the middle column two uncorrected auroral images made by ALIS 19970216 in the auroral green line (5577 Å), and in the right column the same images corrected. Note that the correction is correct for the whole image plane for the image taken with the 70° field-of-view but not for the image with 90° field-of-view where the corners are markedly darker.

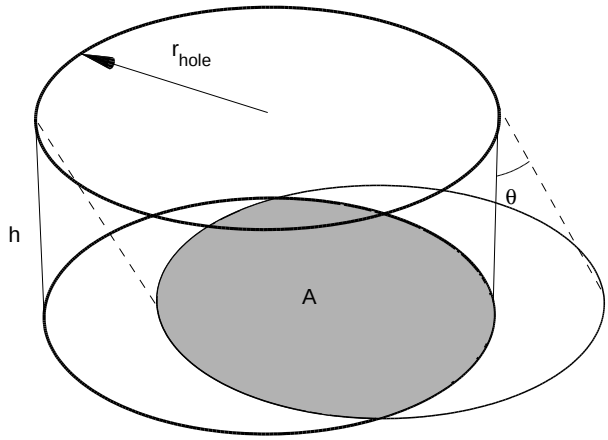

Figure 3.5: Ray-bundles entering a pair of apertures at an angle will be obstructed.

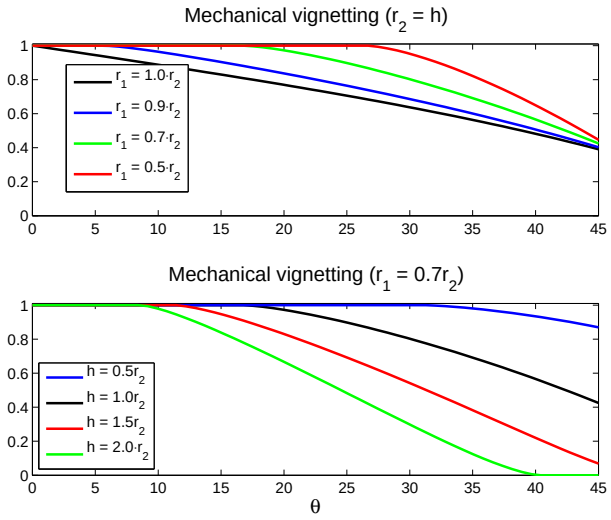

Figure 3.6: The reduction of the clear aperture area light rays can pass through varies with angle of incidence as well as the relation between distance between apertures and aperture diameters.

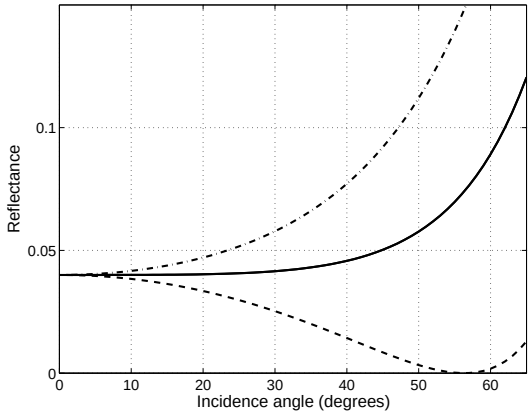

Figure 3.7: Reflectance of transverse and parallel polarised light at an air glass intersection.

the signal, that is already reduced by the optical vignetting equation (3.5), gets significantly weaker than the more central parts of the image and thus is irreparably lost in background noise.

3.5 Transmission of the optics

As a zeroth order assumption, the transmission of the optics is assumed to be constant for all angles θ relative to the optical axis in the full camera field-of-view. For cameras with a narrow field-of-view this can be considered to be a justifiable assumption. The transmission should still be fairly constant for optics with a field-

of-view of about 50°. For fish-eye lenses this remains an slightly open question. However looking at the typical reflectance curves for simple glass (with Brewster angles around 55°) it appears as if it is a good enough approximation out to angles up to some 65°.

3.6 Photo-response non-uniformity

The flat-field correction outlined above take only the large-scale variation into account, not the local pixel-to-pixel sensitivity variation. Still pixels on CCD-arrays have varying sensitivity [e.g. *Holst*, 1998; *Janesick et al.*, 1987; *Burke*, 1996, and references therein] on a pixel to pixel scale, so called photo-response non-uniformity.

Here we will describe a method (based on an idea by P. Rydesäter) to obtain estimates of this from ordinary images obtained during regular observations of the sky. In general the sky brightness varies smoothly on a point-to-point scale and the small scale variations appear randomly distributed both in time and space, except for stars – but they rotate through the image plane when the imager is kept in a fixed rotation. Assuming that intensity gradients vary randomly with a zero average we get an estimate of a correction factor for the pixel-to-pixel variation in sensitivity, C_{p2p}, by taking the average of the ratio between raw image intensity, I_i, and image intensities after 2-D local median filtering, I_i^m, for a large set of images:

$$C_{p2p} = \frac{1}{N} \sum_{i=1}^{N} I_i / I_i^m \tag{3.11}$$

where I_i^m is image I_i 5 by 5 median filtered.

As can be seen C_{p2p} is fairly flat around 1 with 96 percent of the pixels within 1 ± 0.02 and no large scale variations; with some hot and cool pixels. Judging from the small standard deviation of C_{p2p} the estimates seem reliable. The circular pattern that is seen in figure 3.8 is caused by stars that sweep over/through the image plane. Thus their contribution to C_{p2p} is averaged out.

The upper left panel in figure 3.8 shows the estimates of the photo response non-uniformity, PRNU, (point to point sensitivity variation) for ALIS CCD-imager #3. A number of hot and some cool pixels can be seen. In the upper right panel the variance is shown. The circular pattern is caused by stars rotating around Stella Polaris, at pixel coordinate (37,51), and through the pixels. This does not give a large contribution to the sensitivity but contributes to its variance. In the lower left panel is an image before correction for the pixel-to-pixel sensitivity variation, and in the lower right panel the image after correction; and as can be seen a number of the bad pixels have been reasonably corrected.

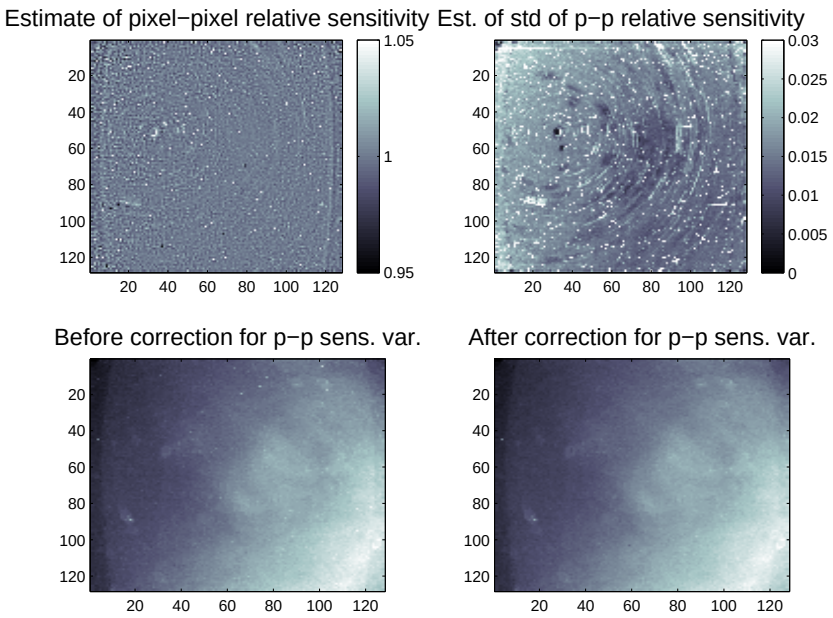

Figure 3.8: The upper left panel shows the estimates of point to point sensitivity variation for ALIS CCD-imager #3. A number of hot and some cool pixels can be seen. In the upper right panel the variance is shown. The circular pattern is caused by stars rotating around Stella Polaris, at pixel coordinate (37,51), and through the pixels. This does not give a large contribution to the sensitivity but contributes to its variance. In the lower left panel is an image before correction for the pixel-to-pixel sensitivity variation, and in the lower right panel the image after correction; and as can be seen a number of the bad pixels have been reasonably corrected.

Chapter 4

Resolution and Aberrations

Imaging systems are not providing unlimited image quality. More introductory stuff goes here...

4.1 Diffraction

If we start by looking at the pin-hole camera we realise that light stemming from a point will not be reproduced as a point in the image plane since the aperture is not a point-hole but a hole with a small diameter d. Thus, light from a point source will make a small spot in the image plane, a so called confusion spot. If we reduce the aperture diameter the confusion spot too reduces in size. But only to a limit, approximately $d = 1.9\sqrt{f\lambda}$, for smaller diameters it again starts to increase - due to diffraction. All apertures in the optics acts as diffraction slits which consequently limits the possible obtainable resolution. The fundamental limit is thus the diffraction limit.

4.2 Aberrations

Practical limits for most "inexpensive" optics, since spherical lenses are not even ideally focusing all rays of light onto a point. Instead there will be some point-spread. There is several causes of this spreading, the third order aberrations (Seidel aberrations) are:

- Coma
- Spherical aberration
- Field curvature
- Astigmatism
- (Distortion) - Circles and lines on circles and lines – and off to complex analysis you go...

and two chromatic:

- Longitudinal
- Transversal

Coma

The first Seidel aberration is coma. This is when there is a tear-shaped point-spread function off the optical axis, either tailing off outwards or inwards. Coma is caused by different refractions by rays hitting the centre and the periphery of the lens. That is rays hitting the lens at all distances from the lens centre is not focused to circles with the same centre in the image plane. The direction of the tail depends on the shape of the lens and as can be seen in figure 4.1 there is a lens shape that removes coma.

Spherical aberration

Spherical aberration is caused by rays entering the lens at different distances from the centre will be focused to different distances. Spherical aberration depends on the shape of the lens.

Field curvature

Typically the focal "plane" is not a plane but rather a curved surface. Figure 4.3 show this for central rays through the same single lens.

Astigmatism

When the focal surface for the central rays are curved, there is typically problems with astigmatism at larger angles from the optical axis. This is because the meridional (red) and sagittal rays have different focal lengths. This makes the shape of the point spread function vary in shape with distance from the nominal focal plane. Since astigmatism is closely related to field curvature this is worsened at large field-of-view angles.

Point-spread-function

Combining diffraction and the different aberrations of the optical system we get a point-spread-function, psf, that may vary in width and shape over the image - then $psf(u, v) \neq psf(u', v')$. If the point spread function of the optics is shift invariant, that is $psf(u, v) \simeq psf(u', v')$ it is possible to move the processing to the Fourier domain where the optical transfer function is defined as: $fft(psf) = otf$. Then all convolutions in the spatial domain (for example $I = I_0 \otimes psf$) turns to multiplications in the Fourier domain ($\tilde{I} = \tilde{I_0} otf$)

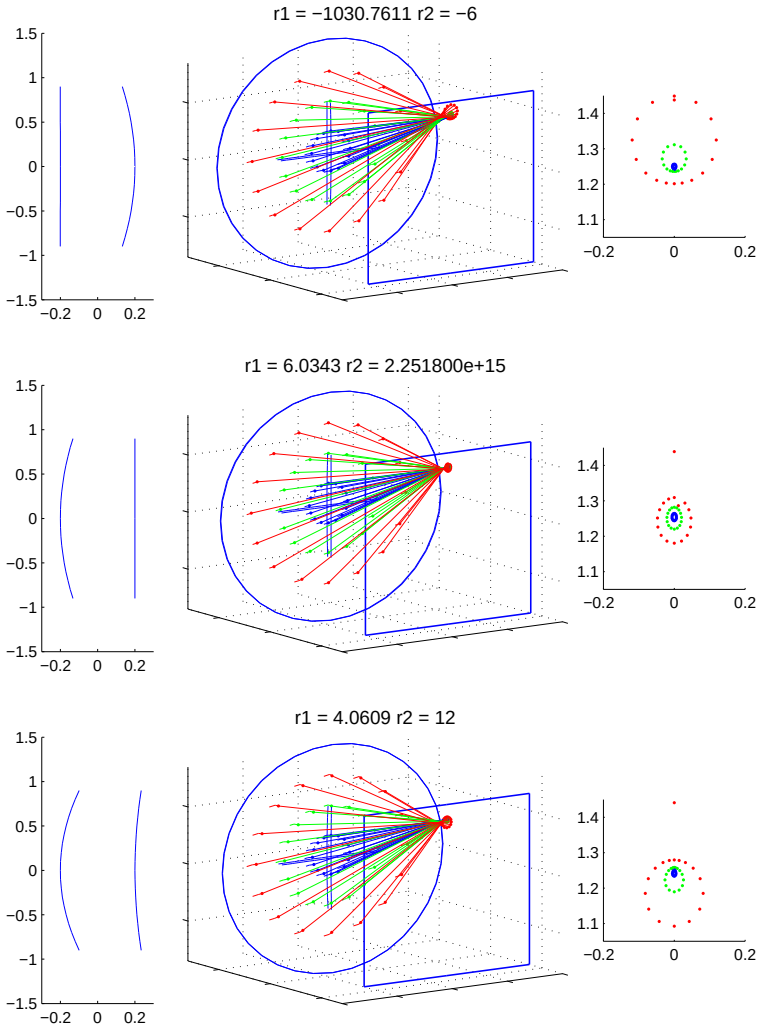

Figure 4.1: Coma is an off axis aberration making points tear shaped either tailing out from or in towards the centre of the image depending on the shape of the lens.

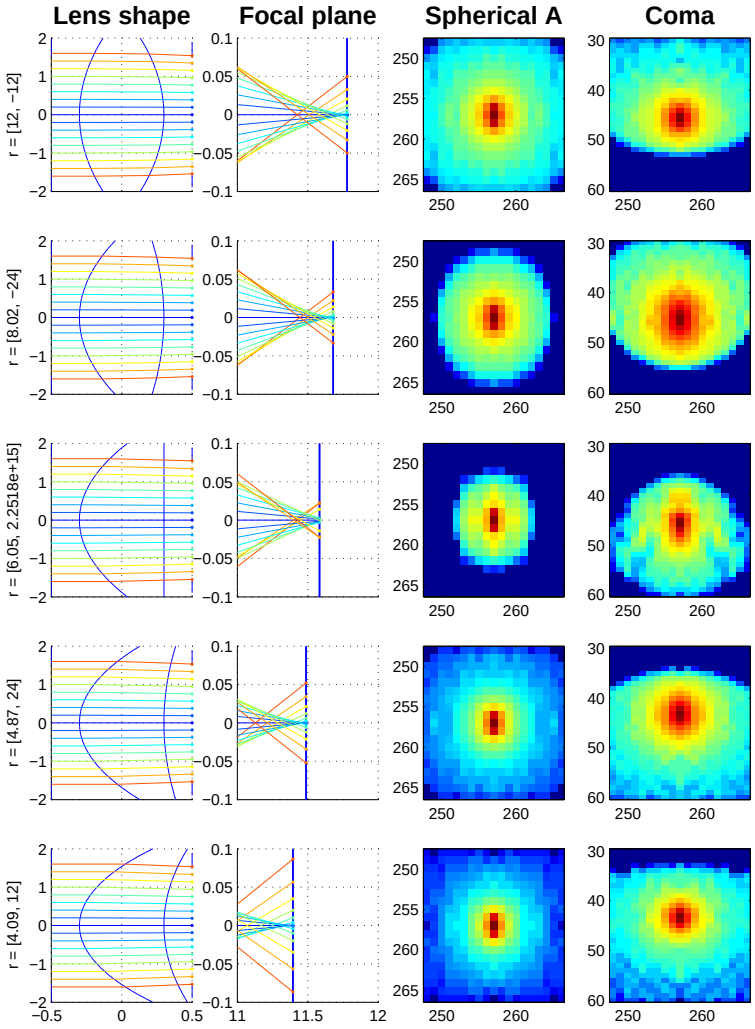

Figure 4.2: Coma and spherical aberration are closely related and both depend on the shape of the lens. Fortunately, for the lens shape that removes coma spherical aberration is close to minimal for a single lens.

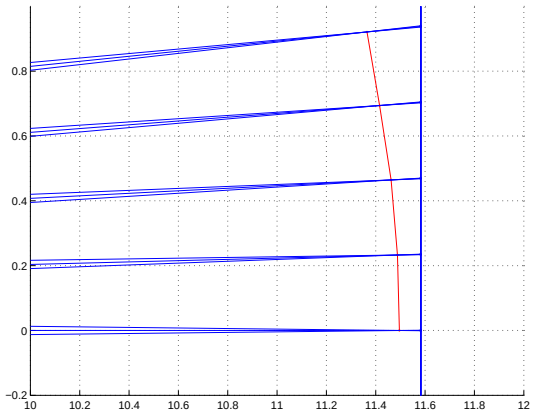

Figure 4.3: Due to spherical lens surfaces the focal "plane" is not free from curvature.

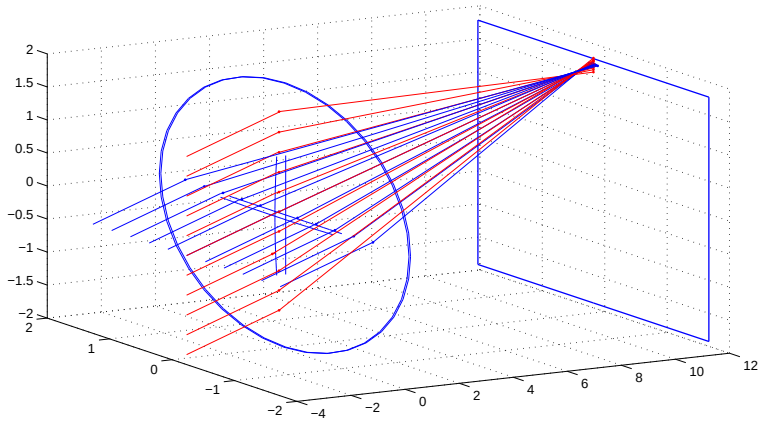

Figure 4.4: Meridional and sagittal rays don't have the same focal distance - this makes the focal "point" non-circular at any image distance.

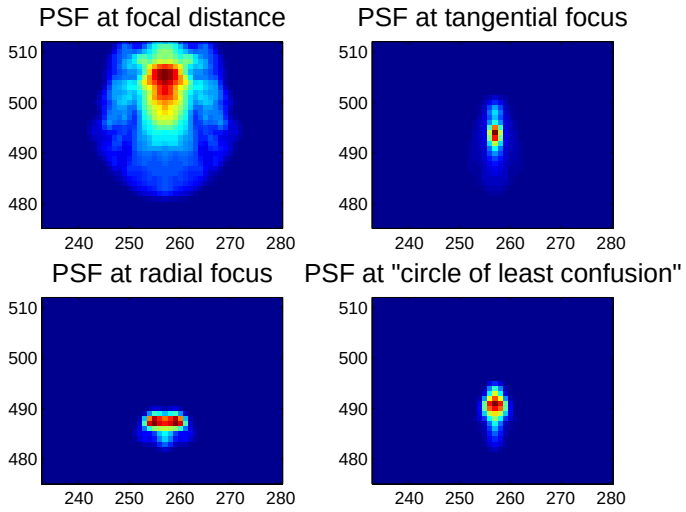

Figure 4.5: Astigmatic aberration makes the meridional and sagittal focus appear at different focal lengths, with a circle of least confusion somewhere in-between.

4.3 Point spread function

For imaging with a medium wide field-of-view, such as in the optics used by ALIS with field-of-view in the range $54° - 90°$, the focus will be aberration limited rather than diffraction limited. For the ALIS cameras this is most likely to be due to astigmatic aberration, as can be seen in Figure 4.6. The effect of this aberration

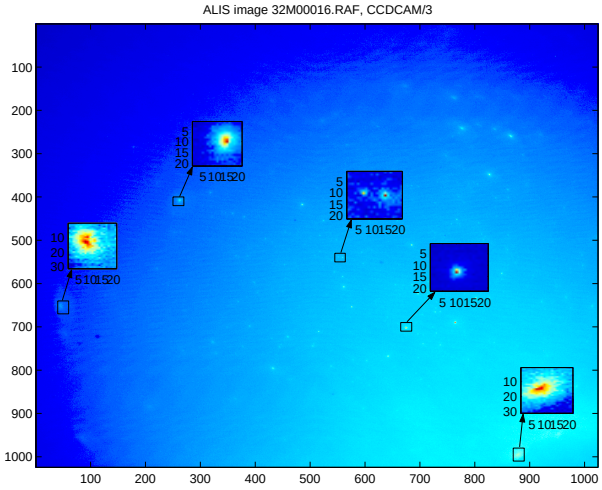

Figure 4.6: Typical ALIS background image with a few selected stars magnified, showing the variation of the PSF with image location.

limited focus is that the point spread function (PSF) varies over the image plate.

For ALIS the PSF can be found by examining the shape of stars in images with long exposure times. From these images it is found that the point spread is fairly well modelled by a two-dimensional Gaussian function with varying widths depending on the location in the image plate. If we take horizontal and vertical slices through the centre of the image, and restrict the study to the variation in PSF width in the radial direction, i.e. neglecting the azimuthal point spread, it is found that the PSF width as a function of distance from the projection of the optical axis varies approximately parabolicly as can be seen in Figure 4.7. This restriction is a "best case" restriction since the increase in the perpendicular width is significantly larger. Such spatial variations in the point-spread obviously violates the requirement of shift invariance required for example for applying deconvolution algorithms based on Fourier transforms.

4.4 Summary

The third-order aberrations varies with the aperture diameter of the lens and the angle relative to the optical axis. In table 4.1 we list the third order dependence of the lens aberrations on the aperture diameter D and distance Y from the image centre. This approximate table might provide useful insight.

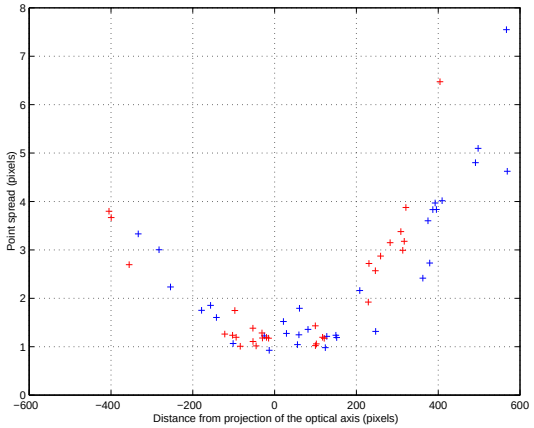

Figure 4.7: Variation in widths of stars in ALIS images taken with an optics with 90° field-of-view.

Spherical aberration	D^3	–
Coma	D^2	Y
Astigmatism	D	Y^2
Field curvature	D	Y^2
Axial colour	D	–
Lateral colour	–	Y

Table 4.1: The variation of Seidel aberrations.

The aberrations are measured by the deviation from the ideal Gaussian image point, either by the size of the associated blur disk on the sensor. For instance, spherical aberration leads to a blur disk for a point source. The diameter of the blur patch increases with the third power of the aperture diameter. Coma only appears at finite distances from the image centre and increases linearly with the distance Y. Photographers can reduce aberrations that depend on the aperture diameter (SA, C, A, FC, AC) by stopping down the lens - but then the speed of the lens is reduced and longer exposures are needed.

Chapter 5

Noise

In which we introduce the villains of the drama.

Noise is an interesting topic. From Wikipedia, the free encyclopedia:

> *In science, and especially in physics and telecommunication, noise is fluctuations in and the addition of external factors to the stream of target information (signal) being received at a detector.*

Thus noise depends on what we are interested in (target information). Before we go off trying to work this out we will benefit from noting that our image intensity can be divided into three parts: the interesting signal, a background intensity and noise. What the interesting signal is of course depends on our focus of interest. For example for an astronomer interested in stellar spectra auroral and airglow emissions are definitely noise and background that affect the dynamic range and the variance in the observations. For an auroral physicist, on the other hand, it is the stars that stick in the eyes. With this notation we separate the noise signal defined above into a part we can remove and one part we can only reduce. The first part might be better denoted deterministic background signal while the second part is random noise.

While the target information is not directly drowned/diluted by a large background signal, the available dynamic range might be severely reduced. And given a relative uncertainty in the calibration of the different background sources, the reduction might not be perfect which then introduces an error.

5.1 Photon counting noise

Since the photon emission process occurs as a random radiative de-excitation process the flux of photons will always have a stochastic nature - thus, the number of photons striking any area at a given time will have a distribution with some width around its expected value.

Poisson - not noise?

From a given number of excited atoms the number of emitted photons will be Poisson distributed - and thankfully the Poisson distribution rapidly enough ($N_{phot} \approx 50$) becomes approximately normal distributed with a probability distribution: $N(N_{phot}, N_{phot}^{1/2})$. For really low count rates the Poissonian characteristics might be important to take into account, otherwise we're good to go with Gaussian distributions and weighted least square methods.

If one should view the random fluctuations in the photon flux as noise is not obvious. The number of photons emitted in a volume is, after all, an observable - and thus, the number of photons actually counted might well be considered true and void of noise.

Still...

...Poisson type noise

The quantum nature of light leads to a randomness in the photon flux. From N excited atoms the number of photons emitted during one second is a Poisson distributed random variable. If M photons is emitted within the field-of-view of a pixel - then the number of photons emitted towards the camera front lens is a binomially distributed random variable.

If we solve the tomographic problem analytically, we will encounter the problem that the forward model is not a deterministic process, but rather a random process

$$\bar{I}_s = \bar{\mathbf{F}}'_{\mathbf{s}} \cdot \bar{\phi} \tag{5.1}$$

where a component ϕ_j of $\bar{\phi}$ is a sample from one random variable with a probability distribution function that depends on several processes from the emission of photons in a volume element j to the detection in the image. Neglecting the randomness in all processes from the camera to the voxel, the random direction of photon emission remains. Since the distribution of emission can safely be assumed to be isotropic and the photon emissions are independent from each other, the probability distribution for the number of photons emitted towards the camera is a binomially distributed random variable:

$$p(\phi_j) = Bi(f_j, p(j, s)) \tag{5.2}$$

where $p(\phi_j)$ is the probability of getting ϕ_j photons, f_j is the total number of photons emitted in voxel j and

$$p(j, s) = A'_{\mathrm{L}}/(4\pi|\bar{r}_{\text{voxel(j)-camera(s)}}|^2) \tag{5.3}$$

is the probability that a photon is emitted towards the front lens of camera s, A'_{L} is the area of the front lens as seen from voxel j. With respect to work on aurora for which the number of photons emitted is large and the probability is small, the binomial distribution in equation (5.2) is well approximated by a normal distribution:

$$Bi(f_j, p(u, v)) \approx N(f_j \cdot p(u, v), \sqrt{f_j \cdot p(u, v)}) \tag{5.4}$$

Thus we see that the imaging process is a stochastic one for which the probability of getting a set of measured images I_s from a source distribution f is:

$$P(\bar{I}_{all}|\bar{f}) = \mathbb{P}(\bar{\mathbf{F}}_{all}, \bar{f}) \tag{5.5}$$

where $P(\bar{I}_{all}|\bar{f})$ is the conditional probability of getting the measured images $\bar{I}_{all}$ given the source distribution $\bar{f}$. For auroral imaging with ALIS the image intensity has a probability distribution that is approximately:

$$p(I) = N(I, \sqrt{I}) \tag{5.6}$$

5.2 Other Noises

In addition to the random variations in the number of photons that create counts in our detectors there are a number of other sources to unwanted signals. This section contains a compact list of some such noise signals with short descriptions about their characteristics.

Cosmic Rays

Cosmic rays occasionally hit the detector and creates a large number of counts in the hit pixels. This is a spatially sparse noise signal with high count-levels in the few random pixels hit. Intensity-wise stars act pretty much the same way as cosmic rays as the intensity from the stars is not part of a Gaussian distribution from aurora or airglow. On the other hand the position of the stars is more predictable.

Large intensities in isolated pixels are relatively easy to remove with median-type filters (as will be described in section 6.3), stars can be selectively removed since we know which pixels are affected.

Interference

Through the wiring high frequency clock signals and also the mains can induce interference patterns on the analog side of the electronics. This if the interference contribution has a stable enough frequency spectrum it is possible to significantly remove this type of noise, in particular if the interference frequencies are large compared to the spectral frequencies with large power in the "true image". If the interference frequencies are relatively low - low meaning that the clean image has considerable spectral power at some of the frequencies with interference reduction becomes more difficult.

Fixed Pattern Noise

The term noise is slightly incorrect for these noise types in that they can be determined/estimated by calibrations. But since they might change with camera setting such as gain and zero level, complete calibration can be difficult to obtain, and they

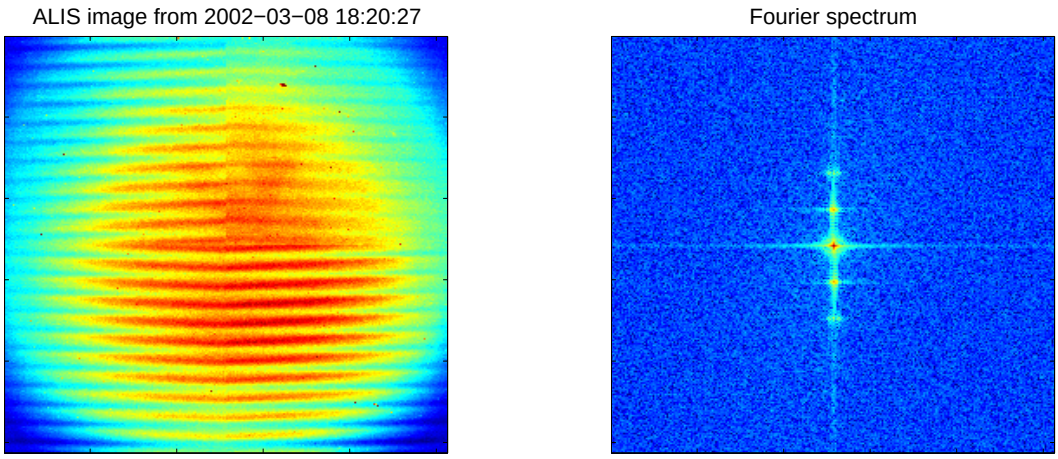

Figure 5.1: Example of an image that are severely affected by interference, to the left, and its Fourier transform, to the right.

might be tricky to eliminate. To be able to completely eliminate these it is preferable to keep the camera settings to a limited set which have carefully calibrated estimates of the bias level, dark current and photo-response non-uniformity. One problem with obtaining accurate estimates of zero-level, dark-current and photo-response non-uniformity images are that these might depend on other parameters such as operating temperature of the sensor, the clock frequencies, and voltages.

Bias level can be determined by zero exposure time images. Dark current can be determined by exposures with lens cover on and different exposure times. In order to obtain accurate estimates it should be kept in mind that the individual zero-exposure images have read noise and that several images should be used to estimate the bias level and its variance.

Pixel to pixel variation in sensitivity can be determined by flat-field imaging (preferably with integrating sphere as light source) or the "statistical flat-fielding" described in section 3.6.

However good these calibrations are there will be some inaccuracies, and it is those that contribute to the truly noisy "fixed pattern noise". The part of the "fixed pattern noise" that is correctly estimated will mainly contribute with reducing the dynamical range available for capturing photons.

Quantisation

When the image intensity is digitised into a low number of levels random variations in the number of photons/counts will be amplified when just on the level between two quantised levels. Please use 12 or 16 bits per pixel. Or make good use of the available grey-levels.

5.3 Noise and resolution

With noise in the images structures that have "intensities at about" the same level as the noise and "limited spatial extend" will be difficult to resolve. If we have a "known" relation that relates a truly ideal image intensities and the noise that appears in an observed image ($I_{obs} = P(I, \bar{p})$) it is interesting to see what structures are statistically significant. For the case when the stochastic relation is as simple as $I_{obs} = N(I, \sqrt{I + C})$, we can do this by a "quad-tree" decomposition (with polynomial least squares fit) of the image until it is impossible to reject the hypothesis that the observed image is plausibly sampled from random variables with the filtered image as the expected values. As can be seen the finer details is irrevocably hidden

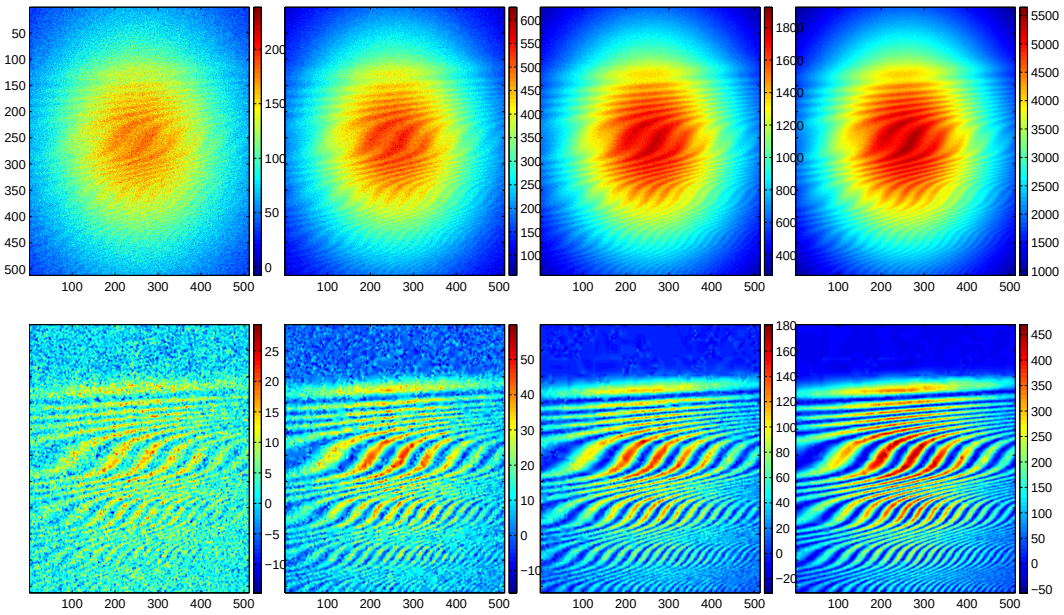

under noise. However, surprisingly much details are statistically significant. This is the fundamental mathematical justification for image filtering, and gives a number of criteria for evaluating image filtering methods. This is dealt with in the following chapter.

Chapter 6

Image filtering

The presence of noise is sometimes a nuisance. When there is a need to remove or reduce the noise we resort to filtering. Filtering can be seen as a science, an art or a dubious hand-craft. However, when the object and the objective is well defined, and careful considerations are made when designing the filter used, filtering is often both successful and useful enough. Provided that uncorrelated noise with expected value close to zero is varying faster than the ideal signal; local averaging makes the contribution from noise decrease faster than the deterioration of the signal. On the other hand, if noise signal with long-range correlation has amplitudes that are confined to narrow enough frequency bands selected Fourier components can be efficiently suppressed without significantly altering the image.

6.1 Statistical Justification for Filtering

With the knowledge of the stochastical characteristics of the imaging process, as outlined in chapter 5, we can look at the filtering process as an attempt to get an estimate of the imaged scene "before the addition of the noise". What we search for is then an image where the pixel intensities are the expected values of the stochastical random process that the observed pixel intensities are sampled from. With images where the relation between intensity and noise is of Poisson-type or $N(I, \sqrt{C_1 I + C_2})$ the standard deviation depends directly on the pixel intensity, the probability distribution of the stochastical process describing the statistical distribution of pixel intensities for any given expected intensity. The saving grace here is that for pixel intensities that have normal distributed probability for the sample intensities it is possible to scale the residual between the observed image, I_o and any suggested expected-intensities image, I_i, to a standard N(0,1) distribution. Still one can create an infinite number of expected-intensities images where the distribution of the intensity residuals follow a normal distribution. The problem here is obviously that we only have one sample per pixel. What we really look for is the simplest image where the pixel intensities give the right distribution of the residuals; with some sensible meaning of "simplest". With the additional assumption that the imaged scene is (piece-wise) smooth "simple" can mean the most smooth image. This leads to the

"total variation" type estimation where an "ideal image" is found by minimising:

$$\min_{I_i(u,v)} : \sum_{u,v} (I_i(u,v) - I_o(u,v))^2 + \lambda \sum_{u,v} |\nabla I_i(u,v)| \tag{6.1}$$

where λ is a smoothing parameter that weights the smoothness of the "ideal" image and the closeness between the "ideal" and observed images. In a slightly different interpretation the "simplest image" can mean the image made with the smallest number of free parameters. With this second approach it is possible to use a number of measures to evaluate both the "simpleness" of the "ideal image" and the fit, or agreement, with the known statistics. One measure often used for accurate fitting of a model to noisy observations is the χ^2 measure:

$$\chi^2 = \frac{1}{N} \sum \frac{(I_i(u,v) - I_o(u,v))^2}{/} \sigma(u,v)^2 \tag{6.2}$$

that should be close to one. This is obviously a necessary condition, but not a sufficient, since χ^2 will be 1 if every pixel $I_i(u,v)$ is equal to $I_o(u,v) + \sigma(u,v)$, or for that matter $I_o(u,v) \pm \sigma(u,v)$ in alternating pixels. This is obviously neither a good fitting image since the residuals have systematic variations, equally important here is that the the residual amplitudes are not normal distributed. One can make a better (even sufficient) measure for the statistical fit by looking at how well the residual distribution fits a normal distribution. Here we will use the squared sum of the difference between the empirical distribution of the residuals and $N(0,1)$, SSdEDN. Other measures balancing model complexity and fit between model and observation are "Bayesian Information Criterion" and "Akaike Information Criterion" (AIC), that both are log-likelihood measures with an additional penalty term that increases with the number of free parameters of the model. AIC is

$$AIC = \sum \frac{(I_i(u,v) - I_o(u,v))^2}{\sigma(u,v)^2} + 2k + \frac{2k(k+1}{n-k-1} \tag{6.3}$$

where k is the number of free parameters of I_i and n is the total number of pixels. We can search for the "ideal image" by fitting two-dimensional least square splines with increasing number of node points to the observed image, and then select the image with that score the best measures. In figure 6.1 one such fitting sequence is presented. In the top panel the noisy synthetic image is shown, and in the second row of panels the best fitting spline-images with 10 by 10, 21 by 21, 32 by 32, 42 by 42 and 53 by 53 node points are shown. There one can see that for the lower number of node points the spline-images are too smooth, with unsatisfying resolution of detail, while for the larger number of node points there start to be some small-scale speckles/ripple. In the third row panel the difference between the scaled residual distribution and a $N(0,1)$ distribution is shown. There one can see that for the lower number of node-points there are too many pixels with too large positive and negative residuals and too few pixels with small residuals; for the larger number of node points the situation is reversed - there are too many pixels with small residuals.

At about 30-by-30 node points the scaled residual distribution is closest to $N(0,1)$. In the bottom panel four measures are presented: A χ^2-measure, plotted as $|\bar{\sigma}-1|$ in blue, the AIC plotted in red, and 2 SSdEDN measures.

This type of "least square fitting spline sequence with statistical measures" filtering make a close to ideal image filter. Unfortunately it is close to prohibitively time-consuming, making it patience-testing in practical use. However, it gives us a means to compare the characteristics of other filters, and some useful measures to evaluate filter performance.

6.2 Linear filtering

If the image intensity varies smoothly and the uncorrelated image noise has a distribution without too wide tails, one can suppress the noise by replacing the pixel intensity with a weighted local average:

$$J(v,u) = \sum_{i=-ws}^{ws} \sum_{j=-wz}^{wz} I(v-j,u-i)w(j,i)$$

Here the size of the filter kernel (w) is 2ws+1 by 2wz+1. There is a plethora of filter kernels for linear filtering. But in image filtering we always use small to very small kernels and there is little to gain by fine tuning w.

Linear filters are efficient against Gaussian and uniform noise, but since it is a linear average it only redistribute and spread out intensity from spikes and salt'n-pepper noise. Further, the local averaging blurs sharp edges and lines with the filter kernel. Thus the filter kernel has to be short compared to the sharpest structures one want to preserve in the image.

matlab notation:

```
J = filter2(w,I,'same');
```

or

```
J = conv2(I,w,'same');
```

or

```
J = imfilter(I,w);
```

Here `I` is the intensity image to be filtered, `w` is the filter kernel. Imfilter has options for handling edges: 'replicate', 'symmetric', or 'circular'.

Filter windows

Since edges and sharp structures are smeared with the filter kernel it is not common to use filter kernels with full-widths at half-maximum larger than 3-7 pixels. A direct consequence of this is that the vast number of filter kernels (with a range of intricate and slightly different pass-band/stop-band/attenuation/and/what-not) in use in signal processing will not differ in any noticeable way, as can be seen in figure 6.3.

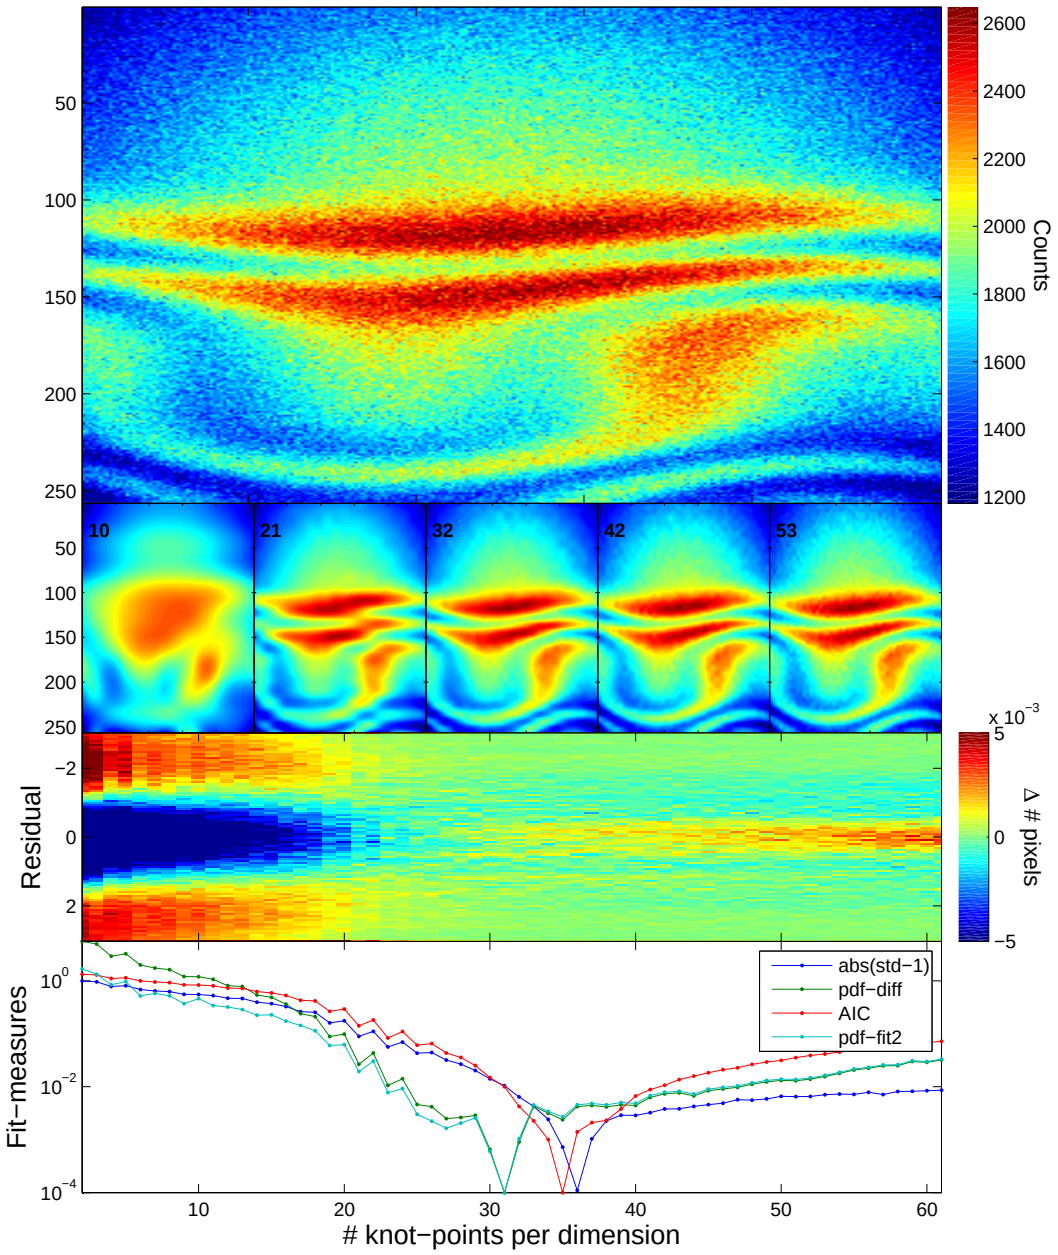

Figure 6.1: .A synthetic noisy image is displayed in the top panel. In the second row of panels a sequence of least-square fitting spline images is shown. By increasing the number of node points of a 2-D least square fitting spline finer and finer details in the image is retrieved.

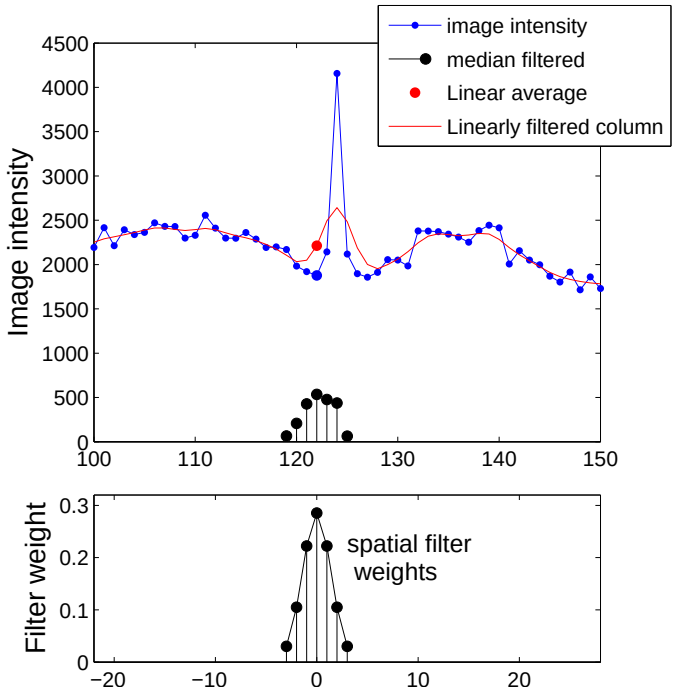

Figure 6.2: In linear filters the image intensity in each pixel, shown in blue in the top panel, is replaced with the average of the neighbouring pixels weighted with the filter kernel, shown in black in the lower panel.

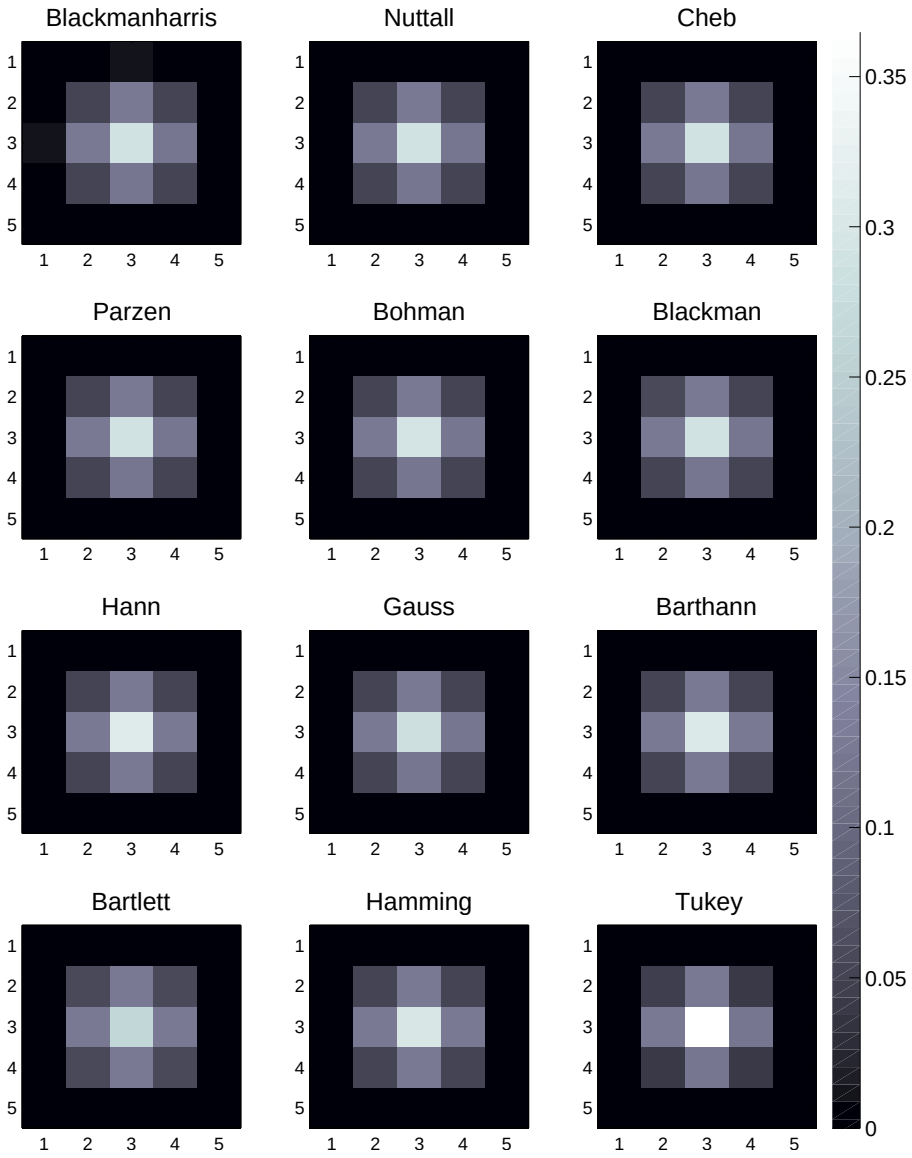

Figure 6.3: Twelve different filter kernels with 3 pixels full-width-at-half-maximum

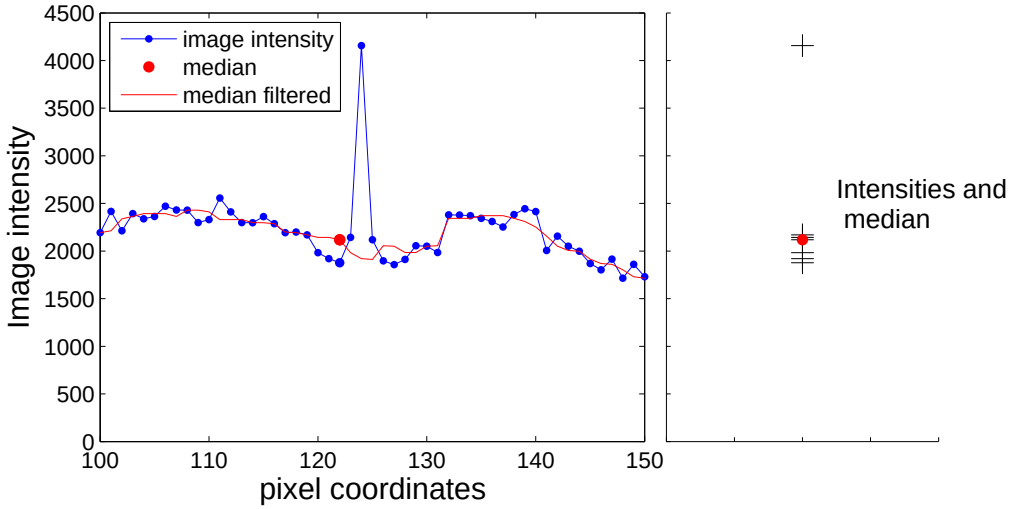

Figure 6.4: In median filters the image intensity in each pixel, shown in blue in the right panel, is replaced with the median intensity, shown as a red dot in both panels, of the neighbouring pixel intensities, shown as black + in the right panel.

(Little linear excursion)

Convolving an image with a Gaussian filter kernel is equivalent to the solution of the diffusion equation in two dimensions

$$\frac{\partial I}{\partial t} = \nabla(\kappa \nabla I) \tag{6.4}$$

with time t from $t = 0$ to $t = t_1$ and initial boundary value $I(t = 0) = I$. When κ is constant the solution to equation (6.4) after a time t_1 is equal to a linear filter with Gaussian kernel with width $\sigma^2 = 2t_1\kappa$.

6.3 Median filters

With median filtering the pixel intensity in a pixel at (u, v) is replaced with the median in the surrounding region:

$$J(v, u) = \text{med}\, I(v - j : v + j, u - i : u + i)$$

Since the impact on the median from the brightest (or faintest) pixels will not depend on its intensity, the median filter is effective against salt and pepper noise. In regions with smooth gradients with low noise, there tends to be a minor stair-case effect. Even worse, if there is periodic noise with wavelengths on the same scale as the filter region, median filter can produce filtered images where the periodic variations are phase shifted.

Matlab notation:

```
J = medfilt2(I,regsize);
```

Here `regsize` is the size $[sx, sy]$ over which to take the sliding 2-D median.

6.4 Sigma-filter

Another way to take the image intensity variation into account when filtering the image can be made with an argument based on the relation between the image noise and the intensity variations. The argument is that given a noise level, n, intensity variations that are small compared to n should be averaged out, while intensity variations that are large can be considered "more statistically significant" and should be preserved. Mathematically this can be obtained by

$$J = \bar{I}_{reg} + \frac{\max(\sigma^2_{reg} - n, 0)}{\max(\sigma^2_{reg}, n)}(I - \bar{I}_{reg}) \tag{6.5}$$

Here $\bar{I}_{reg}$ is the locally averaged image, σ^2_{reg} is the local variance, n is a noise parameter. When $\max(\sigma^2_{reg} - n, 0)$ is much larger than $\max(\sigma^2_{reg}, n)$, the ratio will be close to 1 and the original image intensity will be produced. If on the other hand the noise level is comparable to σ_{reg} the ratio will be close to 0, and the locally averaged will be returned.

This filter does not smear spikes, sharp points and salt'n-pepper noise but in regions with more small-scale intensity variations the noise is averaged out.

matlab notation:

```
J = wiener2(I,regsize);
```

6.5 Susan filter

The SUSAN (Smoothing over Uni-valued Segment Assimilating Nucleus [*Smith and Brady*, 1997], maybe more widely known as "bilateral filter") filter takes the intensity variations into account in a different way than the sigma filter. Instead of weighting with a ratio made up from a noise level and the variance in the surrounding region the filter weights are including a factor that depends of the intensity difference between the central pixel, $I(u, v)$, and the neighbouring pixel $I(v - j, u - i)$.

$$J(v, u) = \frac{\sum_i \sum_j I(v - j, u - i) w(j, i) e^{-(I(v-j,u-i) - I(v,u))^2/\tau^2}}{\sum_i \sum_j w(j, i) e^{-(I(v-j,u-i) - I(v,u))^2/\tau^2}}$$

This way pixels with intensities that differ much from $I(v, u)$ will have a reduced impact on the filtered intensity at $J(v, u)$. This makes the filter very good at preserving edges, corners and fine-scale structures while at the same time reducing the noise level. If the pixel at (v, u) is excluded from the sum susan filters will reduce salt'n-pepper noise and spikes in the data almost as efficiently as median-filters.

matlab notation:

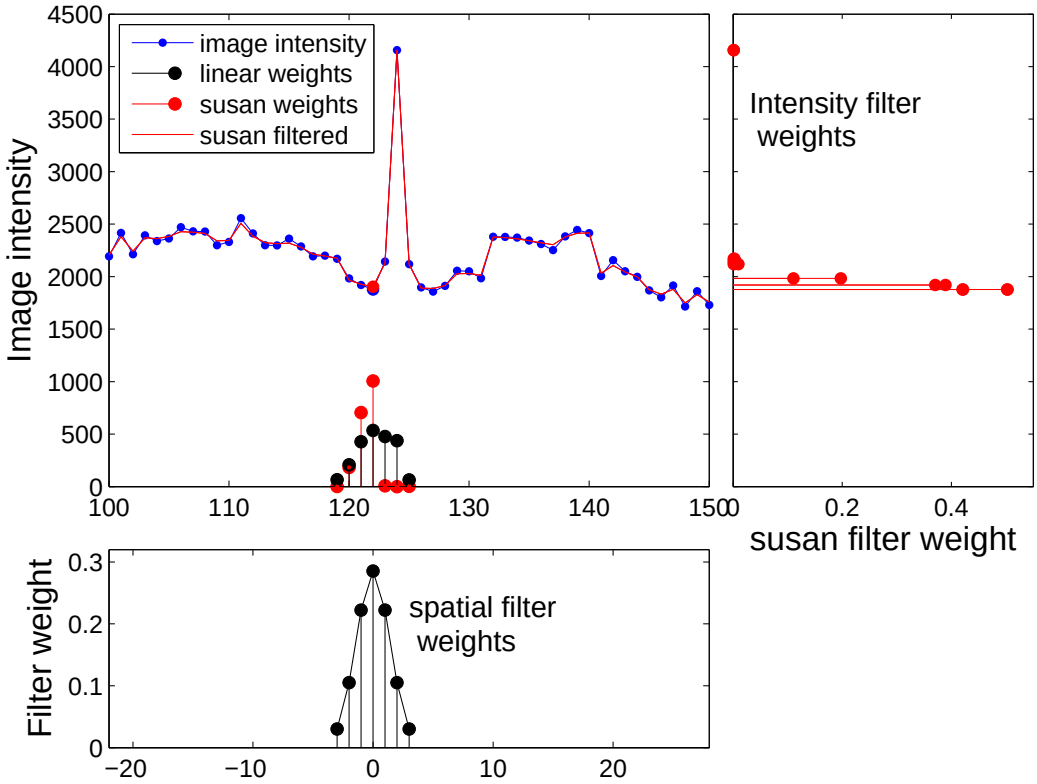

Figure 6.5: For the susan filter the linear filter weights, shown in black in the lower panel, are combined with a factor $e^{-(I(v-j,u-i)-I(v,u))^2/\tau^2}$, shown in the right panel, that reduce the impact of pixels with intensities that deviates much. In the upper left panel the contributions to the intensity of pixel 398 for the susan filter, shown in red is compared to the linear filter, both calculated with the spatial filter kernel shown in the bottom panel. The reduced impact from the pixels with much larger intensities is obvious.

```
OPS = gen_susan;
OPS.tau = max(I.^.5,1000);
OPS.gamma = 2;
OPS.pre_filter = 'n';
OPS.no_center = 1;
J = gen_susan(I,w,OPS);
```

Here `OPS` is an options struct with optional parameters that controls the filtering.

6.6 Non-Linear-Diffusion filters

If we generalise equation (6.4) in such a way that we allow κ do depend on the image intensity, or rather the image gradients we get

$$\frac{\partial I}{\partial t} = \nabla(\kappa(|\nabla I|)\nabla I)$$

If we make the diffusivity function (κ) such that it decrease for steeper image gradients, the flux ($\partial I/\partial t$) can be reduced there. Thus, as the diffusion across steep gradients is inhibited, blurring of steep gradients and edges is reduced.

Several examples of κ have been suggested in the literature. Here we will just mention the diffusivity function used by Weickert (ref!!!)

$$\kappa(g) = 1 - \exp(-C_m/(g/\lambda)^m) \tag{6.6}$$

where m is a shape parameter determining how rapidly the diffusivity decreases when g is larger than λ. C_m is a constant is determined so that the flux ($gd(g)$) is growing for $g < \lambda$ and decreasing for $g > \lambda$. Since the diffusivity decreases exponentially for $|\nabla I|$ above λ the flux decreases as well.

Matlab notation:

```
J = nldif( I, lambda, sigma, m, ...
           stepsize, nosteps);
```

Here `sigma` is the width of the Gaussian kernel to pre-filter the input image I with, before calculating the gradients that determine the local diffusivity.

With this approach much of the image noise can be filtered out while sharp structures can be preserved. One minor problem is that in regions with gradients sharper than the selected cut-off all diffusion is reduced which has the unwanted side-effect of preserving some noise there as well.

A solution to this problem is to make the diffusivity perpendicular to the (Gaussian-smoothed) gradients. Then the noise will be removed by diffusion parallel to the intensity structures.

Matlab notation:

```
J = cedif( I, lambda, sigma, rho, m,...
           stepsize, nosteps, varargin)
```

Here `rho` is the width of the Gaussian kernel to use for calculating the image-structures. Other arguments are as above.

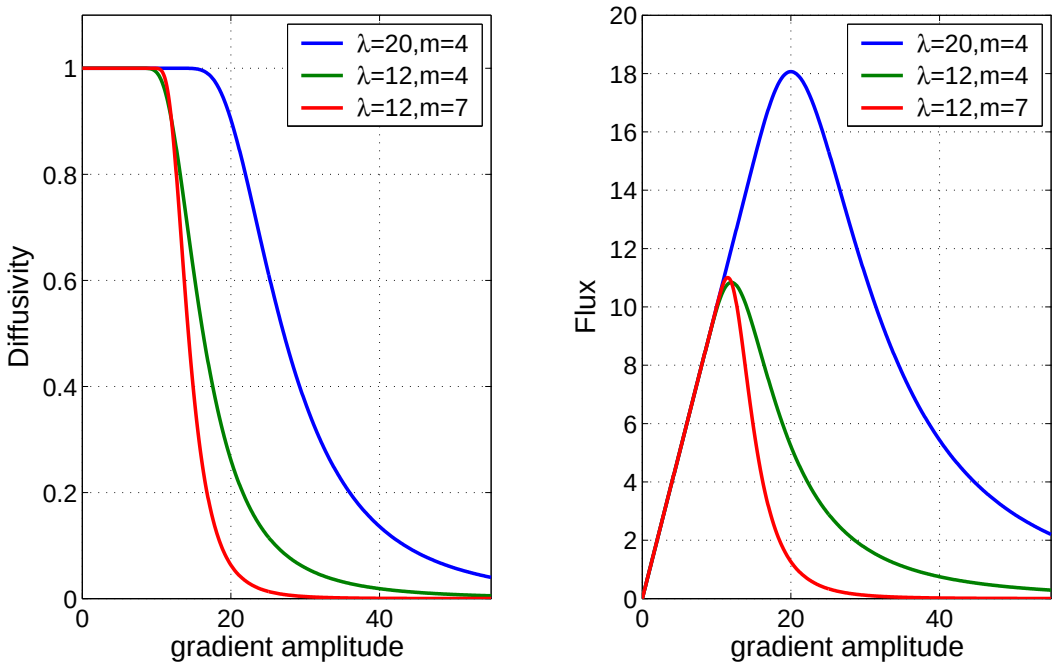

Figure 6.6: Diffusivity and flux functions

6.7 Interference removal

Sometimes there is interference in the images. That is noise at distinct frequencies (not uncorrelated one bit). For this kind of noise the above mentioned low-pass filtering methods will not work. To remove (and prior to removal: to find) the interference spikes we have to do the filtering in the Fourier domain.

Interference and quadrants

Since the ALIS cameras have quad-readout it is not certain that the interference frequencies are the same for all four quadrants, and even if that is the case it is not the case that there is no phase-jump at the quadrant edges. This causes a widening of the interference spikes if we take the Fourier transform over the entire image.

Much better suppression of the interference is obtained if each quadrant is treated separately.

Interference, identification of spikes

In order to remove the interference spike one have to identify which frequency components to remove. This can be done manually, which is tedious but one can make it accurately. Or it can be done automatically. There is one function that identifies and removes these interference frequencies automatically from images:

```
img_out = interference_rem_auto(img_in,...
          if_level,method,wpsm);
```

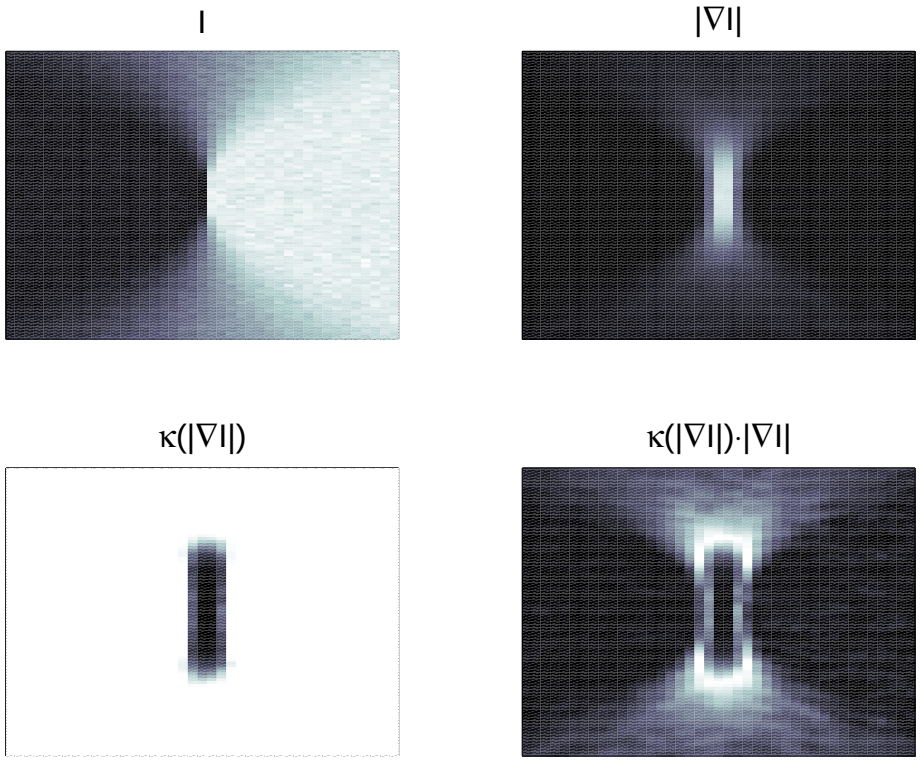

Figure 6.7: An example displaying the workings of the nonlinear diffusion filter parts: upper left displays an example image with the horizontal intensity gradients increasingly steepening towards $[0, 0]$; upper right displays the magnitude of the intensity gradient; lower left displays the diffusivity (according to equation 6.6); lower right displays the nonlinear intensity flux.

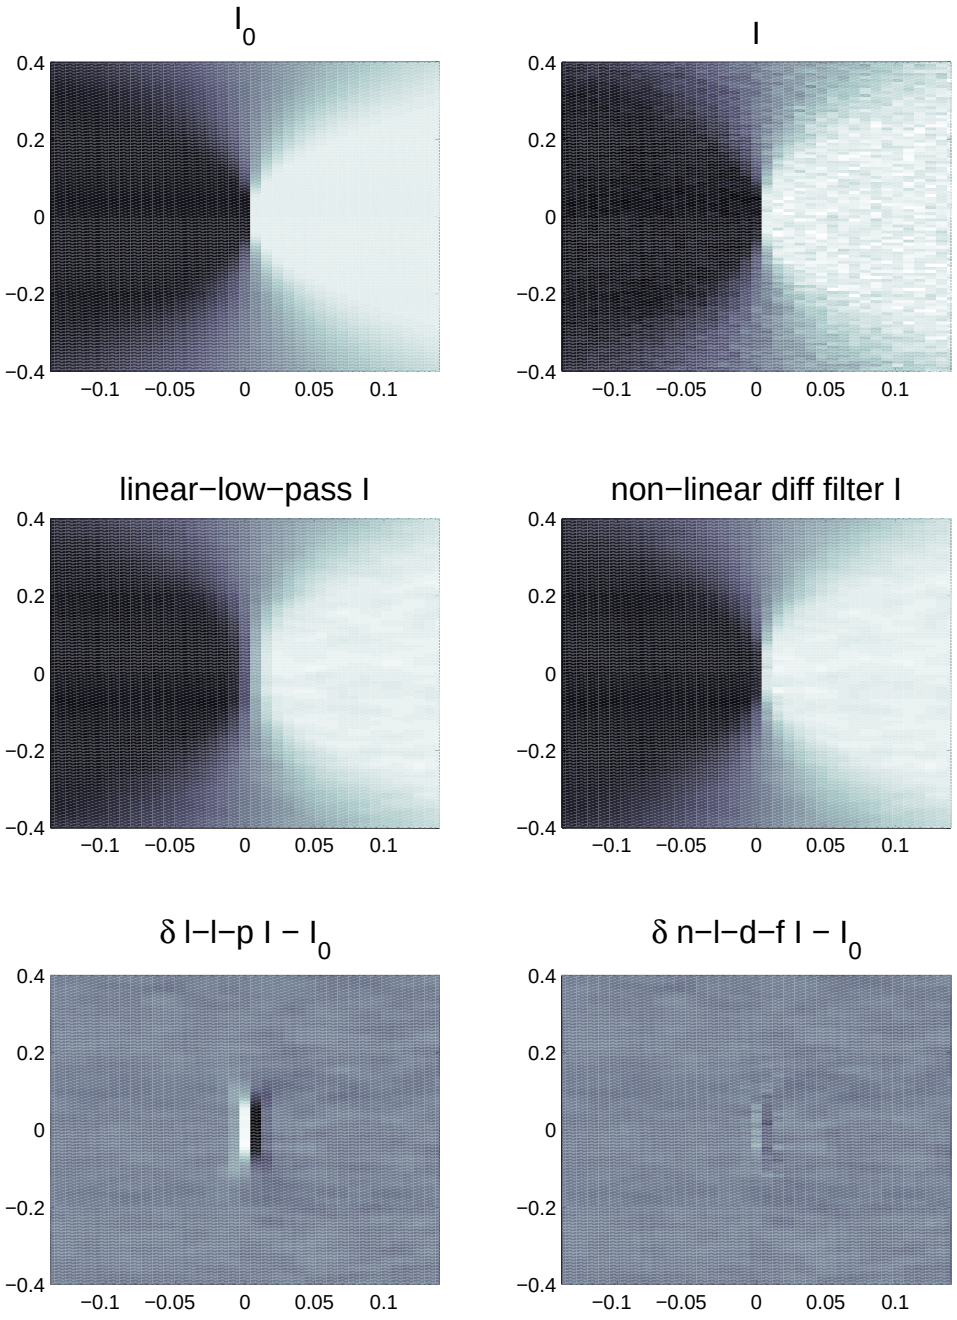

Figure 6.8: An example displaying the difference between linear (diffusion) filtering and nonlinear diffusion filtering. In upper left panel the ideal image ($I = erf(x - y^2)$) is shown; in the upper right panel the ideal image with X percent noise is shown; in the middle panels the linear and nonlinear diffusion filtered images are displayed with the steep gradient in the region around $[0, 0]$ clearly smoothed out by the linear filter in the left panel, which is avoided with the nonlinear diffusion filter shown to the right. In the bottom two panels the difference between the filtered images and the ideal image is displayed, with the linear filtered to the left and the nonlinear diffusion filter to the right.

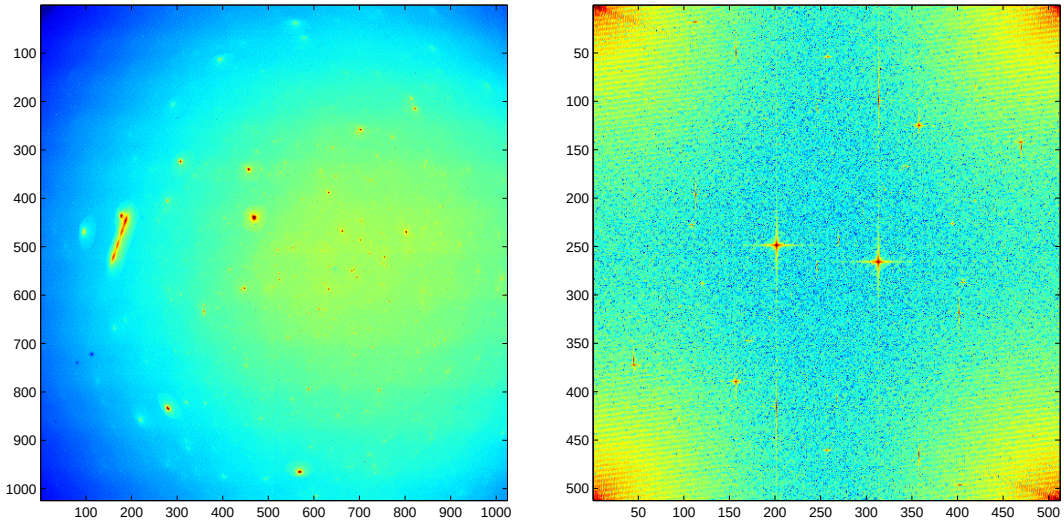

Figure 6.9: An image with coherent noise (interference), shown in the left panel, and its Fourier transform in the right panel, where the frequency components of the interference noise stands out clearly.

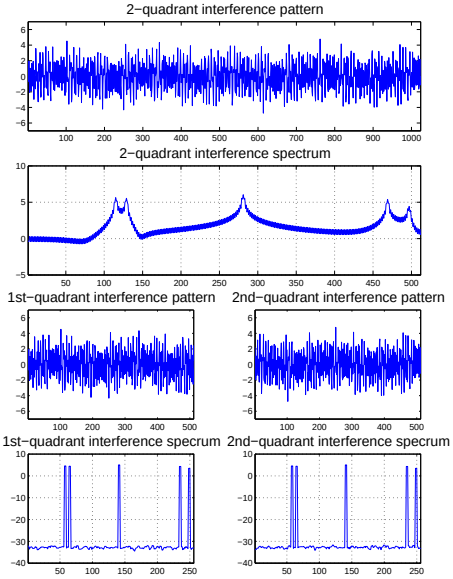

Figure 6.10: For some imagers there are separate read-out electronics for halves or quadrants. When there is interference then there might be phase jumps at the boundary between the different regions. This phase jump leads to a widening of the Fourier spectra if this is not taken into account.

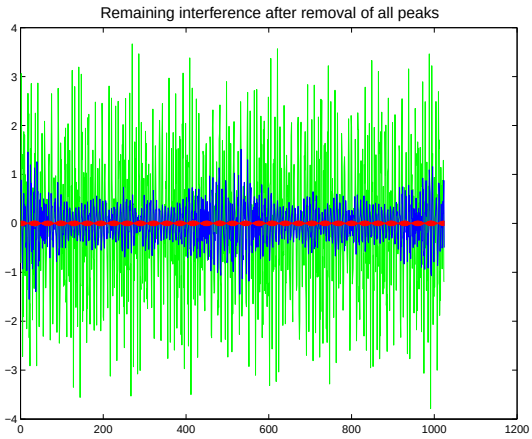

Figure 6.11: The difference in interference suppression for the case of one phase jump at the centre, the green curve is the interference pattern, the blue is with interference suppression without taking into account the phase jump, and the red curve with proper handling of the apse jump.

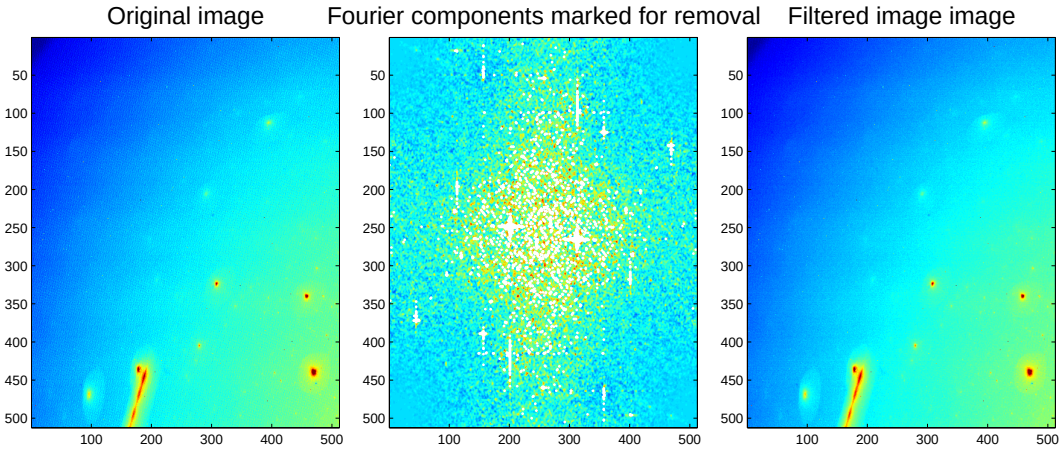

Figure 6.12: For images with high frequency interference, as shown in the left panel, one can identify interference spikes automatically, as shown in the middle panel. The filtered image, shown in the right panel, show the successful noise reduction.

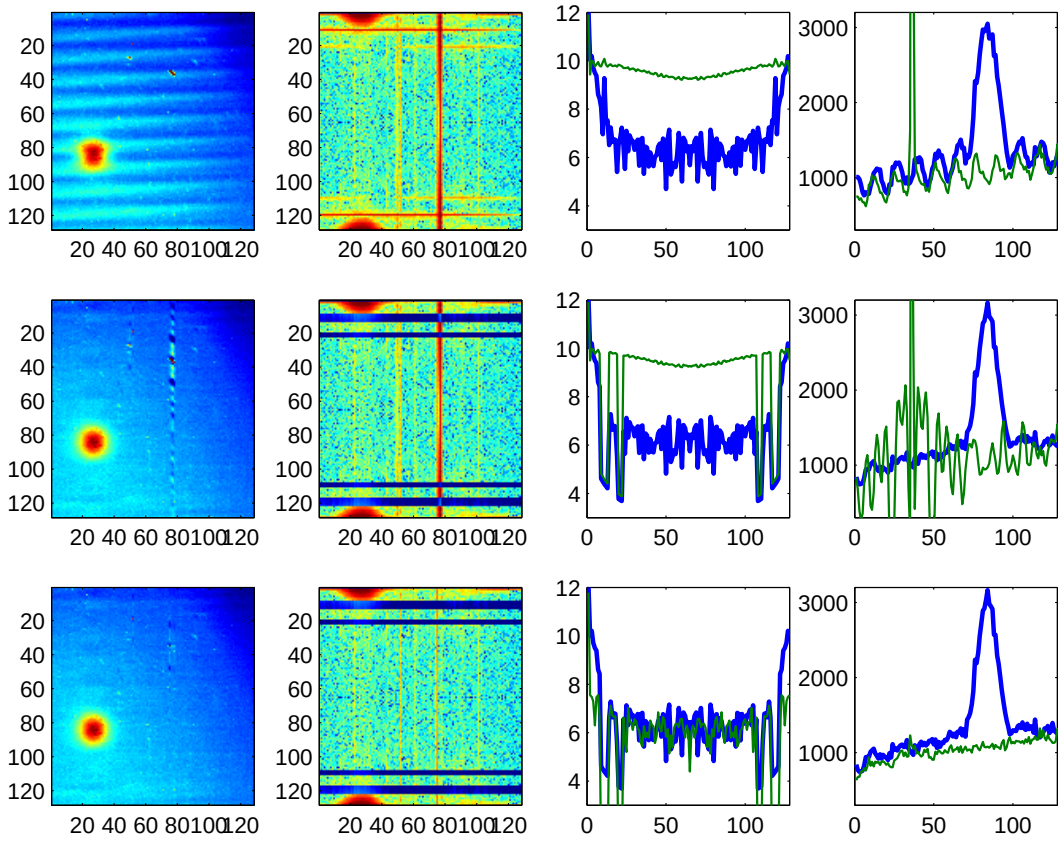

Figure 6.13: When the interference frequency is low the amplitude of the Fourier components from bright point-like sources has to be treated carefully. In the left columns images are shown, in the second column of panels the one dimensional (vertical) Fourier transforms are shown, in the third column from left, two columns of the Fourier transforms are shown, and in the rightmost column of panels the corresponding columns of the images are shown. In the second and third panels in the top row one can see that the bright star has large amplitude Fourier components at all frequencies. In the second row of panels one can see that when removing the Fourier components at the interference frequencies before removing the stellar intensities, there is significant oscillations in the filtered image, this is the inverse of the Gibbs phenomena. In the bottom line of panels the stellar intensity is removed before the interference suppression.

Shown in the left column from top to bottom is the raw image with interference, the image with interference removed before removal of the intensity spikes and the image with the intensity spikes removed before the interference removal. In the second column we show the vertical 1-D Fourier transform of the images int the leftmost column, note the interference bands around line 10 and 20 (and 110 and 120) and the comparatively high amplitudes of all the Fourier components in the columns with the brightest stars - in particular the columns 75–78. In the third column the Fourier spectra of columns 29 (thick/blue) and 77 (thin/green). In the rightmost column the same columns in the images are presented. As can be seen in

the thin/green line the intensity from the star dominates the interference (top row), this leads to oscillating tails when blocking/filtering the interference frequencies (middle row), by first removing the intensity spikes with high amplitude in the Fourier transform at all frequencies this problem can be avoided (bottom row). However, there remains some spurious oscillations from the wider peak (thick/blue line).

6.8 Other filters

There is such a vast number of filters that no-one can plausibly keep track of them all. The filters described in the sections above should cover most needs. The occasion have not yet appeared where I have had to use more specialised filtering techniques, where no combination of the filters above have failed to get the job done. However, to present a few other filters with special applications and some other filters which might be useful to know about - if nothing other just to show that other cunning ways to process data exist - we list a number of other filter methods here:

Harmonic filtering

Harmonic filtering is a slightly modified version of the linear filter

$$J(v,u) = \Big[\sum_{i=-ws}^{ws} \sum_{j=sz}^{wz} \frac{w(i,j)}{I(v-j,u-i)} \Big]^{-1}$$

which has about the same filter characteristics as the linear filters, but is a little better against salt noise.

matlab notation:

```
J = 1./filter2(w,1./I,'same');
```

Homomorphic filters

For images made up by light reflected from objects the image intensity, I(v,u), is a product of the light intensity, L(v,u), and the reflectance of the object, R(v,u):

$$I(v,u) = L(v,u)R(v,u) \tag{6.7}$$

Typically the light (L) varies slowly but much, and the reflectance of the imaged object has much more small-scale variations. When there is need to reduce the impact of the variation in scene illumination one can use the homomorphic filter. The trick with the homomorphic filer is that while one cannot separate the product RL it is trivial to separate the logarithm $\log RL = \log R + \log L$. Then for the cases where the lowest important frequency components of the object has higher frequency than the higher frequency component of the illumination it is possible to reduce the impact of the non-uniform illumination by applying a high-pass filter to the log of the image intensities. As the filter are often to have a low cut-off frequency,

maybe just the first few (5) frequency components, it also more convenient to do the filtering in the Fourier domain.

matlab notation:

```
J = exp(real(ifft2(fft2(log(I)).*...
                   highpassfilter)));
```

Here `highpassfilter` is a high-pass filter in the frequency domain. To preserve the total intensity in the image `highpassfilter(1,1)` should be 1.

- Symmetric nearest neighbour filter (in a (2*n+1)x(2*n+1) region select the pixel with intensity closest to the central pixel of the pair ([i,j],[-i,-j]) then average those.
- α-trimmed mean, take away the $\alpha/2$ brightest and darkest pixels inside the averaging region and replace the central pixel with the average of the remaining pixel intensities.
- Nagao-Matsuyama,
- Kuwahara filter,
- Savitsky-Golay
- ...and on and on...

6.9 Edges and corners

...are problems in image filtering. When we approach edges the region used for local averaging processes start to "spill outside" the image. There are some different (bad) ways to treat this:

- Shortening the filter size.
- Periodic wrapping of image - works for cases where there is reason to assume that the images are periodic. This is what would be implicitly used of filtering was done by FFT-ing the image, multiplying the image fft with the filter-fft and then inverse FFT that product.
- Constant extrapolation/replicating the edge pixels to fill the required number of pixels to cover the filtering window for the edge pixels.
- Symmetric expansion.

Different methods have slightly different effects. My personal preference is to use replication/constant extrapolation, it is simple and stable.

Chapter 7

Image Formation

7.1 Introduction

In this section we will cover the remaining steps that we need to take to make a scene into an image. Put together with the steps outlined in chapters 2 and 3 this will give us a complete forward model that should be sufficient for most modelling and analysis tasks.

For single station imaging it is sufficient to calculate the absolute sensitivity, i.e. the number of photons needed inside the field-of-view of the pixel on the front lens in order to create one count in the image. The primary measurable in imaging is raw counts. With knowledge of the number of photons per count, the effective area of the camera and the pixel field-of-view, it is possible to convert the raw counts to a surface brightness, defined as photons per solid angle per area per time $1/[ster][m^2][s]$. For aeronomic studies it is of more immediate interest to know the total emission from the atmosphere of the spectral features. This conversion from surface brightness is simple provided that the emission is isotropic and not self absorbed. If $\mathbb{N}$ photons have been detected in a pixel with a field-of-view $d\Omega$ and an effective area A_L, an exposure time t and wavelength band $d\lambda$, the corresponding surface brightness is:

$$\mathbb{I} = \frac{\mathbb{N}}{d\Omega A_L d\lambda t} \tag{7.1}$$

This corresponds to the total number of photons emitted from a column with unit area in the direction of the pixel line-of-sight. Assuming that the emission is isotropic, the total number of photons emitted in all directions is:

$$I = 4\pi\mathbb{I} \tag{7.2}$$

In aeronomy the unit for column emission rate has been given the name Rayleigh, $10^6/[m^2][s] = 1[R]$ [*Hunten et al.*, 1956].

For an imaging system intended for tomographic inversion it is necessary to know what fraction of photons emitted in a voxel creates a count in a pixel in the image.

Factors which need to to be taken into account are:

Voxel – pixel field-of-view intersection For all inverse problems in this work there is a voxel representation of the distribution of emission. To calculate the contribution to the image intensity in one pixel from a voxel, the intersection volume, V, between the pixel field-of-view and the voxel must be calculated. This is described in section 7.3.

Pixel field-of-view The pixel field-of-view ($d\Omega$) is solely dependent on the optical characteristics of the camera. A general and straightforward way to calculate this is described in section 3.2.

Atmospheric absorption The light from the aurora and airglow is absorbed in the lower atmosphere, mainly in the stratosphere. This absorption depends on both zenith angle, θ_z, and wavelength, λ. This is described in section 7.9.

Effective collecting area The effective collecting area of the optical system is essentially the size of the front lens as seen from the direction of the voxel. Further the limiting aperture of the optics might change with the angle relative the optical axis. This is described in section 3.2.

Transmission of optics Variation of the transmission of the optical system with angle relative to the optical axis should be accounted for, as described in section 3.5.

Variation in exposure time The exposure time varies slightly from pixel to pixel due to the working of the shutters; a first order correction for this is described in section **??**.

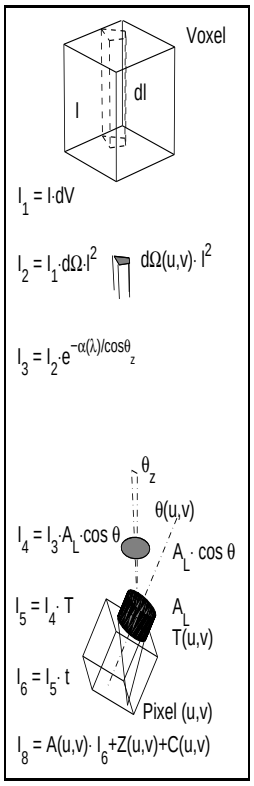

Figure 7.1: Factors in the forward model.

Point spread function The point spread functions (PSF) must be determined. The PSF is the image of a point source. The procedure for determining the PSF is outlined in section 4.3.

Pixel sensitivity The sensitivity of individual pixels must be determined; the necessary requirements are outlined in section 3.6.

7.2 Voxels

When imaging a volume emission, such as aurora, or x-ray attenuation, one have to represent its three-dimensional distribution. This can be done in a number of ways, either as polynomials, Fourier series, splines, or any other set of basis functions. Here we will start with describing voxels (short for volume elements, cf. picture

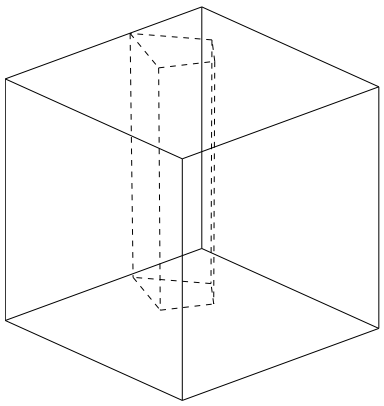

Figure 7.2: The intersection between a pixel field-of-view and a voxel can have different number of faces, here an example with six faces.

element $->$ pixel) that is three-dimensional box-car functions, or square blocks of constant intensity within.

First a few words about voxels.

The simplest set of base functions used in tomographic problems are voxels. It is just a fancy name for a 3-D box with one value for intensity. When voxels are used the number of voxels must be large enough to give a good enough approximation to gradients and curvatures. With to few voxels there will be considerable stair-casing. A good feature of the non-overlapping voxels are that for positive source function ($f(\bar{r}) \geq 0, \text{all } \bar{r}$) the best fit is positive for all voxels ($f_i \geq 0, \text{all } i$).

7.3 Voxel— pixel fov intersection

In order to avoid having to calculate the volume of the complex body with all its bounding planes, the intersection volume dV is approximated by the area spanned by the pixel field-of-view, $|\bar{r}|^2 d\Omega$, at distance, $|\bar{r}|$, to the voxel multiplied by the pixel line-of-sight intersection length dr:

$$dV = |\bar{r}|^2 dr d\Omega \tag{7.3}$$

In order to study the error in the approximation, we turn to the corresponding two-dimensional example illustrated in figure 7.3. The manageable problem is a narrow circle sector intersecting three voxels with intensities $I_0, I_1, \text{and} I_2$, respectively. For a voxel representation of a function in three dimensions to model the function well the function should have slow spatial variations compared to the size of the voxels, i.e. neighbouring voxels should have intensities that do not differ much:

$$I_0 = I_1 + dI_0 \tag{7.4}$$

$$I_2 = I_1 + dI_2 \tag{7.5}$$

Where dI_0 and dI_2 are small compared to I_1.

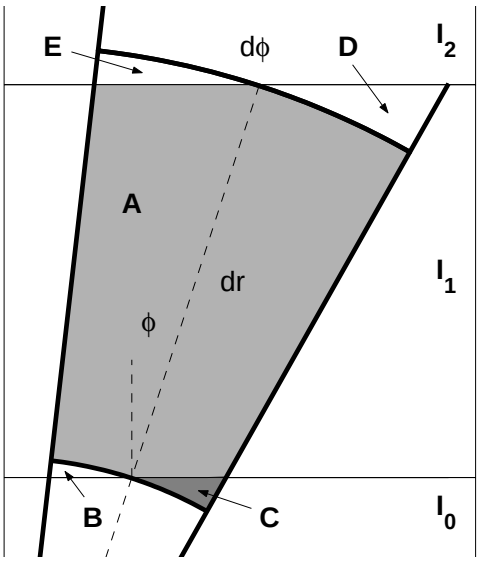

Figure 7.3: Intersection between a "2-D voxel" and a "1-D field-of-view"

The true contribution from voxel 1 is $I_1 (A_A + A_B + A_D)$. The proposed approximation is that the contribution from voxel 1 is given by $I_1 (A_A + A_C + A_E)$. Here A_i is the area of region i. In order to calculate the difference between the exact and the approximate contributions, we calculate the true and approximate contributions from the region $A + B + C + D + E$:

$$I_{true} = A_A \cdot I_1 + A_B \cdot I_1 + A_C \cdot I_0 + A_D \cdot I_1 + A_E \cdot I_2 \tag{7.6}$$

$$I_{app} = A_A \cdot I_1 + A_B \cdot I_0 + A_C \cdot I_1 + A_D \cdot I_2 + A_E \cdot I_1 \tag{7.7}$$

Calculating the difference between the true and the approximate contribution, and using equations (7.4 and 5.5), gives:

$$\Delta I = I_{true} - I_{app} = dI_0 (A_C - A_B) + dI_2 (A_E - A_D) \tag{7.8}$$

After doing some geometry and some algebra the areas of the different regions can be obtained:

$$A_B = r^2 \, d\phi/2 - r^2 \cos^2 \phi \cdot (\tan \phi - \tan(\phi - d\phi/2))/2 \tag{7.9}$$

$$A_C = r^2 \cos^2 \phi \cdot (\tan(\phi + d\phi/2) - \tan \phi)/2 - r^2 \, d\phi/2 \tag{7.10}$$

$$A_D = (r + dr)^2 \cos^2 \phi \cdot (\tan(\phi + d\phi/2) - \tan \phi)/2 - (r + dr)^2 \, d\phi/2 \tag{7.11}$$

$$A_E = (r + dr)^2 \, d\phi/2 - (r + dr)^2 \cos^2 \phi \cdot (\tan \phi - \tan(\phi - d\phi/2))/2 \tag{7.12}$$

Substituting these equations into equation (7.8) and expanding to the third order in $d\phi$, gives:

$$\Delta I = \frac{1}{12} \frac{d\phi^3 \left(2 \cos^2(\phi) - 3\right) \left((dI_2 - dI_0) \, r^2 + 2 \, dI_2 \, r dr + dI_2 \, dr^2\right)}{\cos^2(\phi)} \tag{7.13}$$

Here it is seen that the approximate contribution, I_{app}, is correct to within the third order in $d\phi$ and the second order in (dI, dr). With values typical for ALIS and auroral tomography this gives relative errors that are of the order $5 \cdot 10^{-6}$.

As proven the approximation for voxel–pixel intersection is good enough. A significantly more interesting topic is whether or not the voxel approximation of the continuous volume emission distribution is good enough. As a first test a simple model with only 1 voxel in altitude and a 2-D Gaussian intensity distribution is used with various sizes of the voxels. These voxel representations are projected down onto images with 256 by 256 pixels as shown in figure 7.4. It can be seen

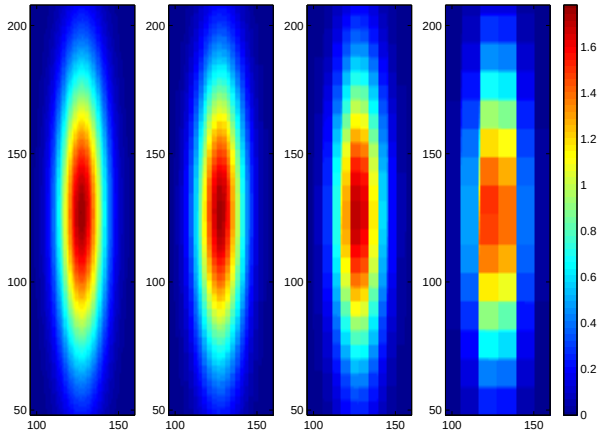

Figure 7.4: Projection of a two-dimensional Gaussian with varying size of the voxels. From left, 256 by 256 voxels with 0.5 km horizontal side length, then respectively 128 by 128 voxels with 1 km side, 64 by 64 voxels with 2 km side and 32 by 32 voxels with 4 km side.

that for voxel sizes that are comparable to the typical dimensions of the structures in the distribution the images become coarse. Looking at the differences between images projected from voxels with different resolutions, as presented in figure 7.5, it is seen that the voxel representation needs to be a posteriori justified, or that the number of voxels must be increased when the result of the tomographic inversion shows structures in the volume emission distribution which have widths of only a few voxels.

7.4 Blobs

There is a viable alternative to voxels for representing three-dimensional distributions. Smooth basis functions of various kinds have been used, with benefits as well as drawbacks. The basic idea with smooth basis functions is that for the same resolution/sizing of basis functions, blobs will approximate smooth three-dimensional distributions better. Problems that arise when using blobs are among others: To avoid perspective effects spherical basis functions are needed; close packed spheres does not fill three-D space completely - the best packing fraction is $\pi/6\sqrt{2}$, for hexag-

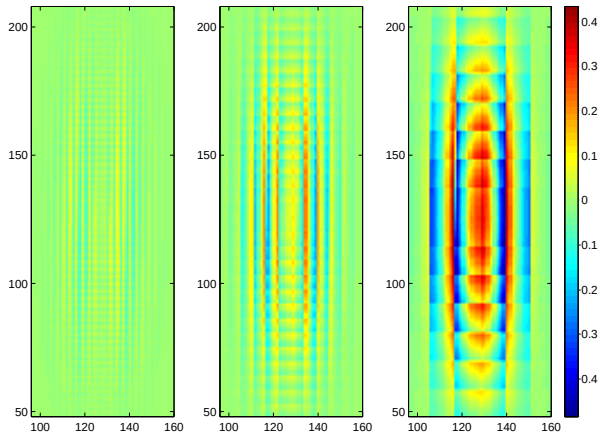

Figure 7.5: Differences between the projection from 256 by 256 voxels and, from left, respectively, projection from 128 by 128 voxels, from 64 by 64 voxels, and from 32 by 32 voxels.

onal close-packed and face centred cubic grids; which are somewhat complex grids. To effectively account for perspective and the variation in size of the basis functions in the image plane one can calculate the foot-print of blobs; then this footprint can be used as a mask (or filter) applied in the image plane for base functions of a given distance and viewing angle; but for a three-D volume which span both some range of distances and viewing directions the number of footprints will become large. To improve the packing fraction one suggestion is to interleave another lattice with smaller blobs displaced relative to the first. This will of course double the number of blobs and increase the computation time for projecting a (actually two) three-D distributions to an image.

Here we will propose the use of "almost spherical" blobs with $\cos^2$ shape. This basis function is completely space filling ($\cos^2(dx_i) + \cos^2(dx_{i+1}) =$ some algebra $= \cos^2 dx_i + \sin^2 dx_i \equiv 1$); further the approximation that $\cos^2(dx_i)\cos^2(dy_i)\cos^2(dz_i)$ is spherical enough to ignore perspective greatly reduces the number of footprints. This actually makes for a fast projection of three-D distributions.

By using smooth overlapping base functions it is possible to reduce stair-casing and get better fits to gradients and sharp curvatures. If the overlap is only to the centre of the nearest neighbours still the positivity requirement is met. To get complete space filling it is not possible to use spherically cylindrical blobs, and to make the projection from blobs to the image plane fast it is necessary for to have blobs that have almost spherically symmetric blobs.

A good compromise is to use base-functions with $\cos^2$ shape in all three directions. This give a complete space filling and negligible perspective effects while at the same time be a smooth representation in the three dimensional space and give smooth projections to the image planes. The perfect space filling taken together with the assumption that neighbouring base functions does not vary to much makes

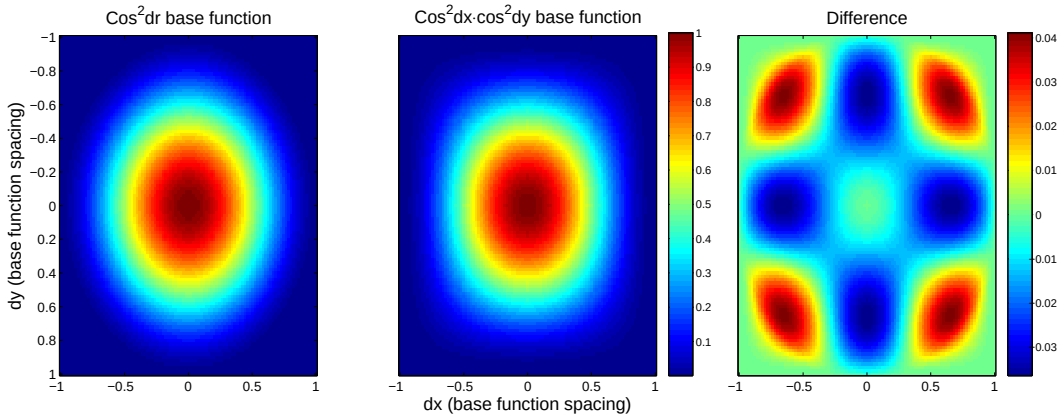

Figure 7.6: Comparison between a circular, $\cos^2 dr$ shown to the left, and non-circular, $\cos^2 dx \cos^2 dy$ shown in the centre, blobs show that the differences, shown at right, is small. This makes the approximation error for line integrals diagonally through the blob small as well.

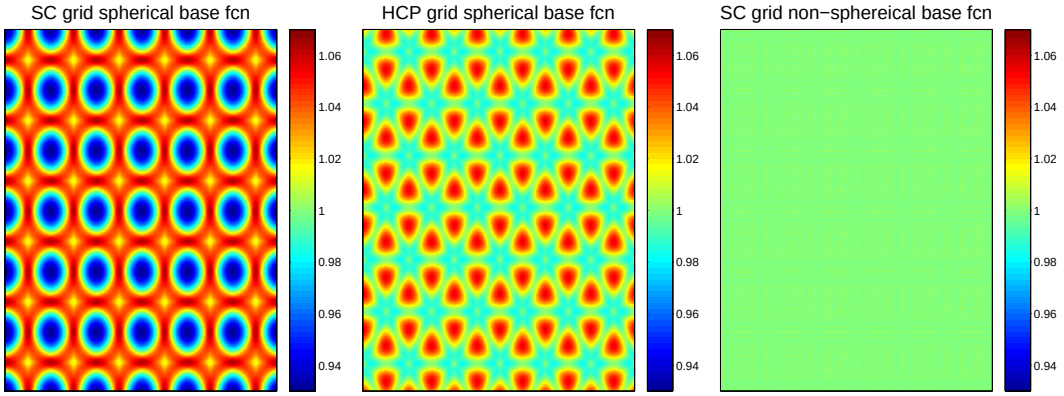

Figure 7.7: Comparisons between spherical blobs filling a simple cubic grid to the left, a hexagonal close packed in centre and a non-spherical blobs on a simple cubic show that only the non-spherical blobs properly fill the space.

the error made in perspective negligible.

7.5 Voxels/Blobs – 1-D comparison

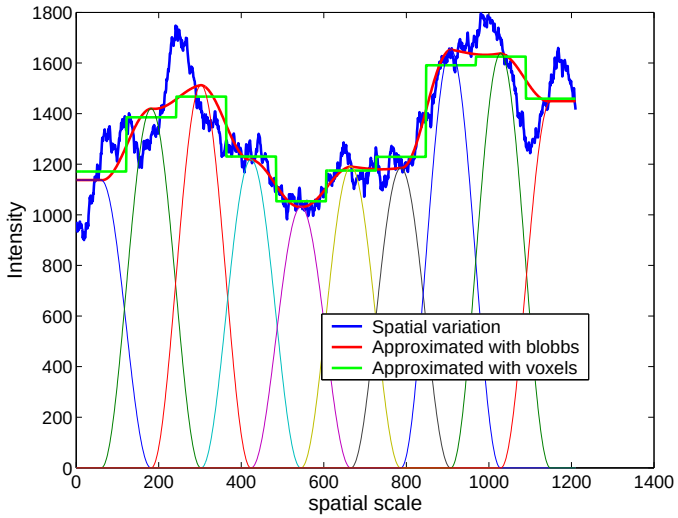

7.6 Fast projection

Further blobs make for really fast approximate (and thanks to its better fit to continuous functions) better projection functions *Rydesäter and Gustavsson* [2001]. This use the fact that if we simplify away the perspective effects we can obtain accurate enough projections by making 2-D filtering in the image plane. This ignores the fact that the blobs are not spherically symmetric which is a small approximation. The projection function then is simplified to:

1. Divide the blobs into size groups depending on their distance from the camera.

2. Project centre points of the blobs to the image plane.

3. Filter the point projections with the corresponding 2-D footprint of the blobs.

4. Repeat step two and three for each group of different sized blobs.

5. Add the separate filtered "group images" together.

6. Possibly multiply with sensitivity factors ($p2p, d\Omega, ...$)

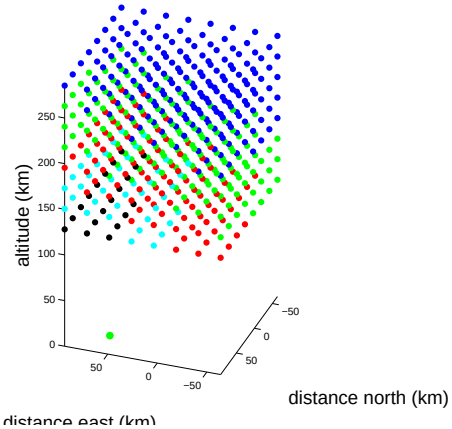

Figure 7.8: The blob are divided into groups based on the distance between them and the camera. Here illustrated with an example where the blobs are divided into five shells. Blobs belonging to different shells are plotted in different colours.

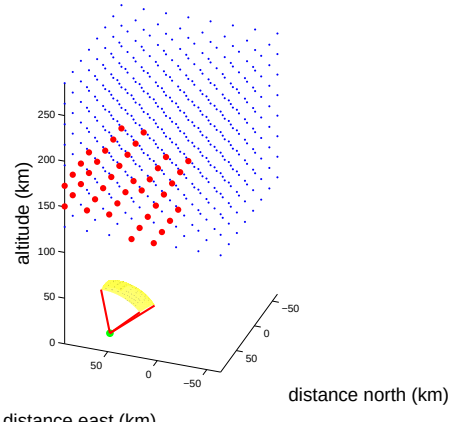

Figure 7.9: For each consecutive layer the image projection of the blob centre positions are calculated, and those pixels are assigned the total image intensity of the blob projection.

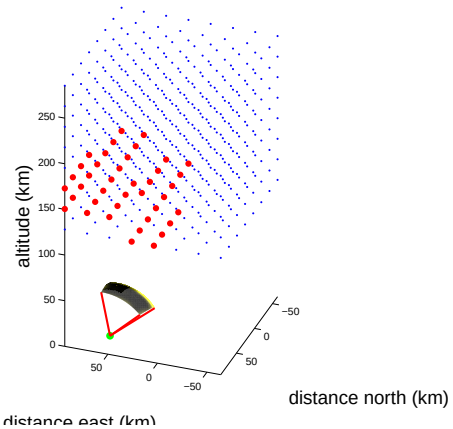

Figure 7.10: Then this image is filtered with a filer-kernel corresponding to the blob footprint.

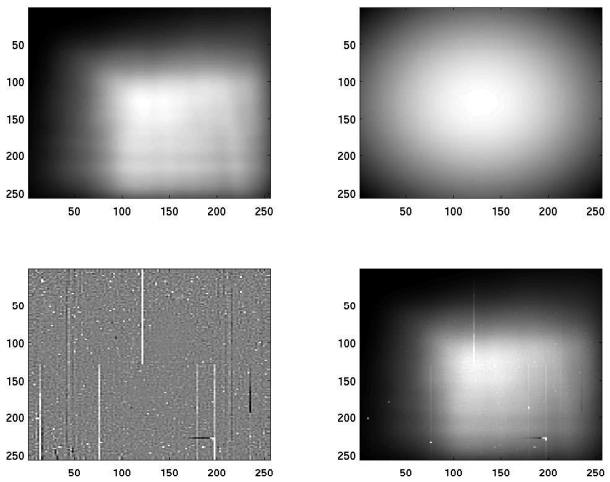

λ (Å)	α
4278	0.236
5577	0.116
6300	0.081
8446	0.080

Table 7.1: Optical depths for the strongest auroral lines

7.7 Voxel–blob comparison

7.8 Voxels and Blobs a summary

Since Blobs gives both better and faster projections there is no apparent reason for using voxels. But since blobs still have no easy implementation where their spatial size varies, there is a last haven for voxels. Provided we work with plaid grids it is no big problem having just about any variation in the size of the voxels.

Otherwise blobs are a superior choice of basis functions.

7.9 Atmospheric absorption

The emissions studied by ALIS are absorbed mainly in the lower atmosphere. Further the emission lines observed by ALIS at $4278, 5577, 6300, 8446$ Å are all optically thin in the ionosphere, i.e. photons emitted in the ionosphere are not absorbed in the region of emission.

The absorption is mainly by the stratospheric ozone. The amount of absorption depends on the length of the light path through the stratosphere. The smallest amount of absorption is obtained for a light path with zero zenith angle. With increasing zenith angle the light path through the stratosphere increases as does the absorption. For moderate zenith angles, $\theta_z < 70°$, the curvature of the earth can be ignored and with a zenith optical depth $\alpha(\lambda)$ that is dependent on wavelength, the absorption is:

$$e^{-\alpha(\lambda)/\cos\theta_z} \tag{7.14}$$

The optical depths by *Arnoldy and Lewis* [1976] that have been used in the current forward model for auroral lines measured by ALIS are given in table 7.1.

Bibliography

Arnoldy, R. J., and P. B. Lewis, Jr, Correlation of ground-based and topside photometric observations with auroral electron spectra measurements at rocket altitudes, *Tech. rep.*, Space Science Center University of New Hampshire, Space Science Center University of New Hampshire Durham, NH 03824, 1976.

Brändström, B. U. E., The Auroral Large Imaging System — Design, operation and scientific results, Ph.D. thesis, Swedish Institute of Space Physics, Kiruna, Sweden, 2003, Sci. Report 279.

Burke, M. W., *Image Acquisition*, Chapman and Hall, 1996.

Glassner, A. S., ed., *Graphics Gems*, Academic Press, 1990, iSBN 0-12-286165-5, 833 pgs.

Gonzalez, R. C., and R. E. Woods, *Digital Image Processing*, Addison-Wesley, 1993.

Green, R. M., *Spherical Astronomy*, Cambridge University Press, 1985.

Hecht, E., and A. Zajac, *Optics*, 2nd ed., Addison-Wesley Series in Physics, Reading, Mass.: Addison-Wesley, 1987, 1987.

Holst, G. C., *CCD arrays, cameras and displays*, 2nd ed., The International Society for Optical Engineering, 1998, ISBN: 0-8194-2853-1.

Hunten, D. M., F. E. Roach, and J. W. Chamberlain, A photometric unit for the airglow and aurora, *J. Atmos. Terr. Phys.*, *8*, 345–346, 1956.

Janesick, J. R., T. Elliott, S. Collins, M. M. Blouke, and J. Freeman, Scientific charge-coupled devices, *Optical Engineering*, *26*, 692–714, 1987.

Rydesäter, P., and B. Gustavsson, Investigation of smooth basis functions and an approximated projection algorithm for faster tomography, *Int. J. Imaging Syst. Technol.*, *11*, 347–354, 2001.

Smart, W. M., *Textbook on Spherical Astronomy*, 6th ed., Cambridge University Press, 1977.

Smith, S. M., and J. M. Brady, SUSAN a new approach to low level image processing, *Int. Journal of Computer Vision*, *23*, 1997.

Index