Detectors and Data Analysis Techniques for Wide Field Optical Imaging

By M.J.Irwin

Royal Greenwich Observatory, Madingley Road, Cambridge, CB3 0EZ, UK

1. Preprocessing CCD Images

"The only uniform CCD is a dead one." – Craig Mackay

Since there are many excellent reviews of CCD data processing (see for example, the reviews by Gullixson (1992) and Gilliland (1992); or the NOAO IRAF photometry cookbook), in this section we will concentrate solely on the steps involved in CCD data pre-processing and highlight some of the problems that may be encountered.

The first stage of any reduction of CCD images involves removing the signature of the CCD from the data. This will involve some or all of the following stages:

- Bias frames – provide both an estimate of the artificial zero-point of the digitised images (to ensure non-negative counts) and an estimate of any fixed pattern noise over the frame. The zero-point is usually monitored throughout the observing session by using the overscan region on each frame to keep track of any variation in bias level, and should generally vary by less than 1%. The "trailing" overscan region is generated by continuing to clock out the CCD after all the charge has been readout and as such solely measures the intrinsic bias level and *rms* noise of the electronics. Note that there is usually an equivalent "leading" underscan region, whose main purpose is to "ramp" up the electronics before real data is read out. Consequently, the overscan region should provide a more reliable estimate of the bias level in the frame. For a modern CCD the bias frame should be a constant, flat level, ie. characterisable by a single number. For an example of unmistakable pattern noise, see figure 9 of Hanisch (1992). If a two-dimensional bias correction is necessary, multiple bias frames need to be taken and combined to remove cosmic ray hits. All other frames are then bias-corrected and overscan regions, underscan regions and any other unwanted strips trimmed off.

- Dark current – the dark current on modern CCDs is usually $\sim$1 e^- per hour and can be ignored. Otherwise, long exposure dark frames are needed to characterise the dark current. As noted earlier, the dark current is both strongly temperature sensitive and power on/off sensitive. The dark frames need to be stacked to remove cosmic-rays and then subtracted off all bias-corrected data frames. The amount to be subtracted is proportional to the relative exposure times. A further advantage of taking a series of dark frames is that the cosmic ray sensitivity can be characterised – useful for optimising cosmic-ray removal routines.

- Flat fields – measure the two-dimensional pixel sensitivity map, including large scale gradients and vignetting caused by dust particles on the filters or CCD window etc.. There are three basic approaches: take dome flats by attempting to uniformly illuminate either the interior of the dome or a white screen using a tungsten light source – hard to get uniform illumination but good for pixel-to-pixel sensitivity measurements; use the twilight sky – problems are short window of opportunity and sky illumination gradients can cause problems for wide field CCDs; make a "superflat" from the astronomical targets using either a shift and add technique, or a series of deep exposures on different regions

– not always feasible. Whichever method is used the frames have to be first combined, to remove either cosmic rays or stars/galaxies, before being used to normalise out the sensitivity variations.

- Fringing – an awkward problem with both wavelength and spatial dependence. Only affects thinned CCDs toward red limit of optical spectrum (see previous section). Sky fringing has the added difficulty of being time-varying too, since the relative intensities of the narrow night sky emission lines (ie. those that cause the sky fringing) varies with atmospheric conditions. Continuum sources (eg. most astronomical objects and tungsten dome lamps) are not seriously affected. Note particularly that twilight sky flats will show fringing in the R, I and Z passbands and as such are not suitable for flatfielding the CCD. If fringing is a problem and the dome flats suffer from non-uniform illumination then use the dome flats to remove only the pixel-to-pixel sensitivity variations and use the sky flats to remove the large scale gradients (although even these may be affected by fringing). From the point-of-view of the astronomical objects, the sky fringing is additive, and hence is best estimated and removed by forming a "supersky" by stacking offset target exposures taken under similar atmospheric conditions.

- Cosmic rays – in target frames the effect can be reduced in three generic ways: the optimum method is to modify subsequent analysis programs to ignore them; if this is impractical cosmic-ray removers based on single frames work by varaints on the theme of looking for isolated pixels with intensities/gradients "well above" the local background/gradient – helped by knowledge of the probability distribution function (PDF) of cosmic ray hits (from dark frames) and a good noise model for the detector (readout noise and photon noise); better is to use multiple frames and perform a pixel-by-pixel comparison using the noise model as previously and some form of k-sigma clipping to reject cosmic rays – see next section.

- Other problems – assorted and include hot spots, traps, dead pixels/columns, saturation, video pattern noise, CTE and smearing, deferred charge etc.. Refer to Gullixson (1992) and Gilliland (1992) and references therein for more details on recognising and dealing with them.

To be a little more precise, we can summarise the various stages of CCD data preprocessing in a compact manner borrowing heavily from the formalism of Gullixson (1992). The raw CCD frame is related to the previous list of corrections by,

$$raw = (obj + sky \times (1 + fringe)) \times qe + dark + flash - skim + bias \tag{1.1}$$

where most of the terms are as above, apart from $flash$, which denotes any pre-flash (low level illumination of CCD using on-board LEDs) that may be applied prior to exposing the CCD in order to reduce the non-linearities caused by deferred charge transfer/skimming, $skim$.

We first correct for the Zero Exposure Additive Spatial Systematics by forming average frames viz;

$$\begin{aligned} < dark > &= dark + flash - skim + bias \\ < skim > &= flash - skim + bias \\ < bias > &= bias \end{aligned}$$

and use the overscan region to monitor and remove variations in bias DC level so that,

$$raw - < .... > -(ovscn - < ovscn >) = (obj + sky \times (1 + fringe)) \times qe \tag{1.2}$$

where $< .... >$ denotes whichever of the 3 corrections or combinations thereof is appropriate or necessary.

To correct for Multiplicative Spatial Systematics, ie. flat fielding, form

$$< flat >= qe + dark + flash - skim + bias \tag{1.3}$$

$$const \times \frac{raw - < .... >}{< flat > - < .... >} = obj + sky \times (1 + fringe) \tag{1.4}$$

where *const* is a normalising constant chosen to make the average multiplicative qe correction unity.

If fringing is present, or if for some reason the twilight sky flat does not result in "flat" frames, correct for Residual Spatial Systematics – shift and stare and stack, to form

$$< sky >= sky \times (1 + fringe); \quad < sky >= sky \tag{1.5}$$

and either additively apply, if fringing is the cause, or multiplicatively apply if problem was non-uniform flats.

Finally, although this may be necessary in the preceding steps as well, register the individual frames, including geometric distortion corrections as necessary, using variants of the general 6-plate constant linear transformation,

$$x' = a.x + b.y + c; \quad y' = d.x + e.y + f \tag{1.6}$$

and combine. It is simplest to compute the linear transforms by matching the coordinates of pairs of images from each frame. Note that if only shifts are required, $a = 1;\ b = 0;\ d = 0;\ e = 1$; if only a small angle of rotation and shift are required, $a = 1;\ b = -d;\ e = 1$.

1.1. *Combining Images*

A common theme in the previous section for both removing the effects of unwanted sources (stars, cosmic-rays etc..) and to improve the signal-to-noise of images, involves combining image frames. Is there an optimum way to do this ? If the frames are aligned astrometrically and scaled photometrically, such that pixels on different frames can be directly combined then the problem is well defined. Otherwise, suitable frame-to-frame transformations as in equation 1.6 need to be applied and possibly some sort of intensity mapping of the form,

$$new = raw \times scale + shift \tag{1.7}$$

to correct for varying sky levels or atmospheric transmission.

Let us now consider a series of matched and scaled values of a particular pixel, with intensities in each of m frames, x_i, for $i = 1, m$. For any underlying PDF of the intensities, $P(x)$, which may include contributions from photon noise, readout noise, cosmic rays, bad pixels etc.., various estimators of central location, ie. the "average" value, can be defined. For example:

- the MEAN minimises

$$< (x_i - \hat{x})^2 >_i = \int (x - \hat{x})^2 P(x).dx$$

- the MEDIAN minimises

$$< |x_i - \hat{x}| >_i = \int |x - \hat{x}| P(x).dx$$

- the MODE, or a local fit to the MODE, can be used to estimate the position of the maximum of $P(x)$ but requires an impractically large number of frames to be available.

• If, and only if, $P(x)$ is exactly Gaussian (or Poisson), say with variance σ^2, is the MEAN the optimum estimator. In which case the error in the "average" values, $\hat{x}$ would be $\sigma/\sqrt{m-1}$. For such idealised cases, the MEDIAN would have error $\sqrt{\pi/2} \times \sigma/\sqrt{m-1}$, but for practical cases would have the added advantage of not being as susceptible to outlying points.

In practice, perfect Gaussian/Poisson distributions are rarely encountered so how do we generalise estimators of central location to cope with "real noise" ?

Tukey and others (eg. Hoaglin, Mosteller & Tukey 1983) have pioneered the development of so-called robust statistics, or m-estimators, to cope with exactly this problem. An m-estimator of central location can be thought of as a way of creating more robust Maximum Likelihood Estimators (MLEs) applicable to "real" noise distributions. Two of the best estimators appear to be Tukey's biweight function and Andrew's wave function. To illustrate the concept, the m-estimators are defined to be the solution of

$$\sum_i \psi(u_i) = \sum_i w(u_i).u_i = 0 \tag{1.8}$$

where $w(u_i)$ can be thought of as a weighting function; u_i is a measure of the relative deviation of the data from the solution

$$u_i = \frac{x_i - T}{c.S} \tag{1.9}$$

with T the best estimate, S the scatter (ie. a measure of the *rms* spread) and c a tuning constant. In practice equation 1.8 is solved iteratively using

$$\hat{T} = \frac{\sum_i w(u_i).x_i}{\sum_i w(u_i)} \tag{1.10}$$

which at convergence is equivalent to solving equation 1.8 directly. Tukey's biweight function is defined by

$$\begin{aligned} \psi(u) &= u.(1-u^2)^2, \quad |u| \le 1 \\ &= 0, \quad |u| > 1 \end{aligned} \tag{1.11}$$

and Andrew's wave function by

$$\begin{aligned} \psi(u) &= \frac{1}{\pi} sin(\pi.u), \quad |u| \le 1 \\ &= 0, \quad |u| > 1 \end{aligned} \tag{1.12}$$

with $S = MAD$ (ie. the Median of the Absolute Deviation from the median, a popular robust estimator of sigma, equivalent to $\sigma_{rms}/0.6745$ for a Gaussian distribution) and c set to between $\approx$6 to 12. Note that both are equivalent to clipping outlying members (cf. k-sigma clipping) from contributing and both perform better than the straight median estimator, whilst remaining robust to aberrant points. For comparison, a k-sigma clipped mean estimator would be specified by $\psi(u) = u, \quad |u| \le 1; \quad = 0, \quad |u| > 1$.

The condition satisified by equation 1.8 arises naturally from minimising a function of the form

$$\sum_i \rho(u_i); \qquad \frac{d\rho(u)}{du} = \psi(u) \tag{1.13}$$

and suggests a simple equivalence with generalised MLEs. For example: consider a process whereby $x_i = \overline{x} + \epsilon_i$ and the PDF of ϵ_i is given by

$$P(\epsilon_i) = \alpha.N(0, \sigma_i) + \beta.U(-a, a) \tag{1.14}$$

where $N(0, \sigma_i)$ represents the Gaussian core of the noise distribution, and the uniform distribution, $U(-a, a)$, represents the non-Gaussian extended tail. If the range $[-a, a]$ is large compared to the core size and the fraction of outlying points is low, $\beta << 1$, then to a good approximation this distribution is equivalent to

$$\begin{aligned} P(\epsilon_i) &= \frac{\alpha}{\sqrt{2\pi\sigma_i^2}}.exp(\frac{-\epsilon_i^2}{2\sigma_i^2}), \quad |\epsilon_i| \leq k\sigma_i \\ &= \beta, \quad k\sigma_i < |\epsilon_i| \leq a \\ &= 0, \quad |\epsilon_i| > a \end{aligned} \tag{1.15}$$

or, we could equally well have postulated this form for the PDF at the beginning. Assuming independent measurements, x_i, minimising the log-likelihood function then leads to the following estimator for $\hat{x}$,

$$\hat{x} = \frac{\sum_{i=1}^{m'} x_i/\sigma_i^2}{\sum_{i=1}^{m'} 1/\sigma_i^2} \tag{1.16}$$

where m' denotes the observations within the k-sigma clipped range. Equation 1.16 has to be solved iteratively since the clipping boundary is a function of the current estimates of $\hat{x}$ (and $\hat{\sigma}^2$). The well known k-sigma clipping is therefore equivalent to a MLE with a flat extended error tail and illustrates the general applicability of the MLE approach. Note further that equation 1.16 without the k-sigma clipping restriction, is simply the conventional inverse variance weighting scheme, usually derived directly from the requirement that the combined image signal:to:noise ratio should be a maximum. If a more realistic PDF for the non-Gaussian/ Poisson errors is known, including what fraction of pixels are likely to be affected (eg. from the PDF of cosmic rays hits on dark frames), then it is worthwhile considering a numerical solution to the MLE problem that makes direct use of this information.

2. Detecting and Parameterising Objects

The previous section dealt with some of the rationale behind removing the detector-dependent signatures from two-dimensional CCD images. We now turn to the problems encountered in detecting and parameterising any objects that may be present in the two-dimensional data. Since, the signature of the detector defines what device was used to create the image, once that signature has been removed, subsequent analysis techniques should be (more-or-less) device independent, ie. applicable to CCD data, digitised photographic data, infra-red array detectors, and so on.

In order to break down the overall analysis strategy into manageable portions the following main tasks can be identified (see Irwin (1985) and references therein for further discussion on these and related points):

(*a*) Estimate the local sky background over the field and track any variations at adequate resolution to eventually remove them.

(*b*) Detect objects/blends of objects and keep a list of pixels belonging to each blend for further analysis.

(*c*) Parameterise the detected objects, ie. perform astrometry, photometry and some sort of shape analysis.

In deciding on a rationale for these three stages, natural questions to consider are: is it feasible to do the whole analysis automatically – after all wide field optical astronomy has the potential to generate Gbytes of data per night; how does a human observer interactively interpret two-dimensional maps and what can we learn from this; how accurate are any derived image parameters likely to be; and closely related to this last question, which methods of analysis will give the best results at reasonable computational cost ?

2.1. *Sky Background Analysis*

To track any potential background variation, the whole field can be partitioned up into an array of suitably-sized "background regions". (Here, suitable means, a compromise between small enough size to follow real variations and large enough size to reduce subsequent error in the estimate.) For each region compute the pixel intensity histogram and derive various statistics from it, as illustrated in figure 1. A straight mean is clearly hopeless, since even in regions free of faint objects, the odd pixel defect can be enough to seriously bias it. The median is a significant improvement, but owing to the innate skewness of the histogram will always tend to overestimate the local sky level. A popular method, and one which I recommend, is to fit a Gaussian to the core (ie. within say ± 1 sigma) of the histogram, as a robust estimator of the mode, an esimate of sigma is a useful byproduct. This works well in relatively uncrowded regions but again systematically overestimates the sky level as the image number density goes up (see Irwin (1985) for examples). Finally, modelling the skewness of the histogram explicitly (Bijaoui, 1980) can in principle give the most unbiassed answer, although unfortunately the strong correlation between the asymmetry in the data (usually a free-parameter in the model) and the sky location can lead to instability problems.

A problem with histogram methods, particularly with full 16-bit data, is selecting a sensible bin size – too fine a sampling and the mode is ill-defined; too coarse and resolution is lost due to quantisation. A useful rule of thumb is to select a bin size such that the noise level at sky, $\sigma_{sky} \approx 10$ units. Gaussian fitting to the modal core region can readily recover information at a 1/10th bin size level. This is usually much smaller than either the systematic errors caused by contaminating objects, or the random error in the sky estimate $\approx N_{pixels}^{-1/2} \times \sigma_{sky}$; where N_{pixels} is the number of pixels in the histogram within $\pm 1\sigma_{sky}$ of the sky level.

Combining all the background region sky estimates then gives a two-dimensional array of "background" values that can be further processed to help remove regions affected by, for example, the halos of bright stars or galaxies. A local sky value at any point in the original field can then be obtained by interpolation from the background array. This is the approach adopted by the APM facility in automatically processing photographic plates (see Kibblewhite *et al.* , 1984 and Irwin & Trimble, 1984). Note that the process of estimating the background, or sky level, has been completely decoupled from estimating the parameters of an object.

The preceding approach can be thought of as a variant on the scheme of digitally filtering the full two-dimensional pixel map to remove "slowly varying" features, ie. background variations. A popular alternative based on this analogy, is to use something like a two-dimensional median filter (Tukey, 1971) followed by a linear filter (to reduce the step-like appearance median filtering leaves behind), to track the slowly varying features. These are then subtracted off to leave a sky-less map. This method itself is of course a variant of unsharp-masking (eg. Pratt, 1978).

A final alternative is to estimate the background on an image-by-image basis (eg. Newell, 1982). To do this, define an "annulus" around each image, with inner radius several psf radii away, and an outer radius as large as possible. Within the annulus compute either a k-sigma clipped mean estimate, or use the modal estimator $mode = 3 \times median - 2 \times mean$. The latter formula is a good working approximation for a generalised Gaussian distribution of the form N(0,1) $\times(1 + \alpha.(x^3 - 3x))$, but I would recommend the clipped mean because of its generality.

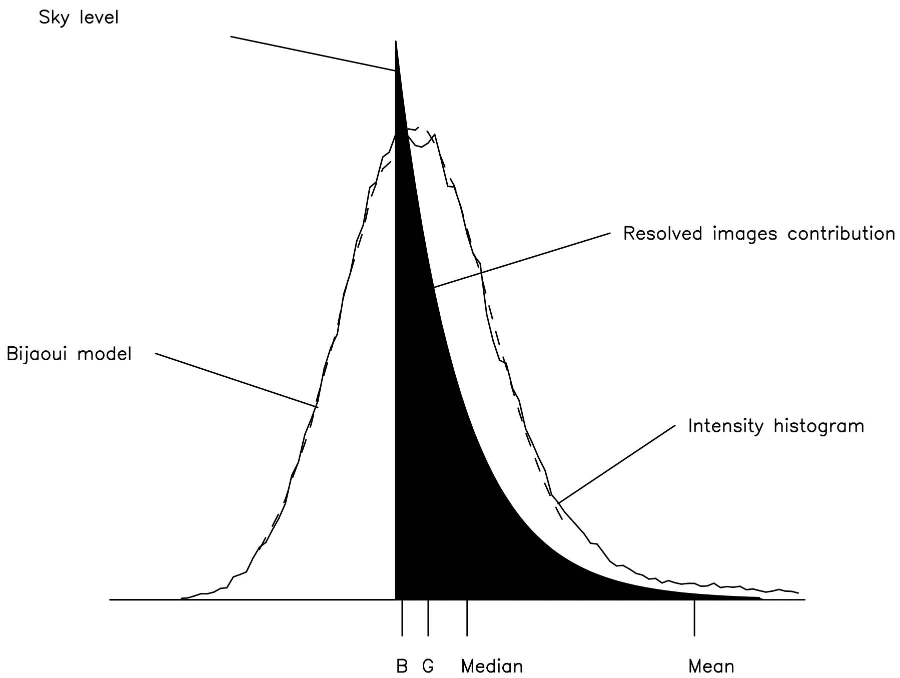

FIGURE 1. Sky background estimation from assorted statistics of the pixel intensity histogram. The histogram of the shaded region is convolved with the PDF of the noise to give the observed intensity distribution. Note in particular how the most widely used estimators of sky level are ordered: Mean > Median > mode (from Gaussian fit - G) > Bijaoui model fit (B), due to the skewness caused by resolved images. A popular variant, the k-sigma clipped mean computed directly from the data, can be as good as the Gaussian fit to the mode.

There are three main problems with the image-by-image method. First, setting appropriate annulii boundaries when dealing with mixtures of stars and galaxies of unknown profile (plus the problems of neighbouring images), is non-trivial. Second, although the combination of spatial clipping, inherent in the annulus approach, and intensity clipping seems like an improvement on the histogram, remember that the psf of star images is much larger than a few core radii and it is precisely these same low level, but large area, intensities near sky, that cause the trouble with either the k-sigma clipping algorithm or the histogram method. A final criticism, is that while a histogram-based method attempts to make full use of all the pixel information and the continuity of the background values, this is not true of the local image method.

2.2. *Optimal Detection of Images*

The classical approach for optimum detection of signals (images) in the presence of random noise is to apply a matched detection filter to the raw data prior to setting detection criteria (see for example Pratt, 1978). To a good approximation the random noise present on the detector comes from a Gaussian/Poisson process and is easy to deal with. However, the presence of other images in the data, particularly faint stars and background galaxies, and possible systematic artefacts due to the detector, complicates things by introducing a "systematic" noise component. Here we define any artefact/image outside of the image of interest, to be a potential systematic noise problem. A good example of this arises in the problem of detecting Low Surface Brightness Galaxies (LSBGs). Since LSBGs often have scale sizes of ≈1 arcmin, multitudes of other faint images will generally be superimposed and make the identification task more complex.

Woodward & Davies (1958) presented a useful insight to the problem of optimal detection of images using a straightforward application of Bayes' theorem. The argument runs as follows: consider the question of detecting the reflected echo of a radar pulse of known shape, $x_{t+\tau}$, but unknown lag τ, buried within a noisy, discretely sampled, time series, y_t, such that

$$y_t = x_{t+\tau} + \epsilon_t \tag{2.17}$$

where ϵ_t is assumed to be random Gaussian noise of variance σ_t^2 †. The aims are to locate the signal and in doing so accurately estimate the position, τ. We first construct the likelihood function of observing the sequences, $x_t, y_t;\ t = 1, ...N$, for a particular value of τ,

$$L(\mathbf{x}, \mathbf{y} \mid \tau) = \prod_{t=1}^{N} P(\epsilon_t \mid \tau) \tag{2.18}$$

and then use Bayes' theorem to argue that in the absence of prior information, finding the maximum of the likelihood function is equivalent to finding the most probable value of τ given $\mathbf{x}$ and $\mathbf{y}$. Now

$$P(\tau \mid \mathbf{x}, \mathbf{y}) \propto exp(-\sum_{t=1}^{N} \frac{\epsilon_t^2}{2\sigma_t^2}) \tag{2.19}$$

and maximising the probability is equivalent to a weighted least-squares minimisation of

$$\sum_t \frac{\epsilon_t^2}{\sigma_t^2} = \sum_t \frac{1}{\sigma_t^2}.(y_t - x_{t+\tau})^2 \tag{2.20}$$

or in the case of $\sigma_t^2 = constant$, finding the maximum of the cross- correlation of y_t with $x_{t+\tau}$. Not only does this illustrate how a cross-correlation function is related to the solution we seek, it shows how to generalise the traditional cross-correlation approach to deal with more complex forms of noise contamination.

To go back to our original problem of detecting images in noise; since the only images difficult to locate are those marginally above the sky noise, a constant noise approximation is a valid assumption and certainly for a CCD, the noise in adjacent pixels is to all intents and purposes random. In other words, to optimally detect images, cross-correlate ‡ the entire frame with the appropriate filter; the psf in the case of point-like sources and possibly an exponential function of chosen scale length for larger faint galaxies.

From a practical viewpoint a matched filter maximises the resultant peak signal:to:*rms* noise. The same criteria can indeed be used to derive it. Figure 2 illustrates this improvement for a series of Gaussian target objects using both Gaussian and top-hat detection filters of various sizes. Defining the two-dimensional target profile as $I = I_p.exp(-r^2/\sigma^2)$ and the normalised detection filter as $I = exp(-r^2/\sigma_g^2)/\pi\sigma_g^2$ changes the peak intensity, $I_p \rightarrow I_p \div (1 + \frac{\sigma_g^2}{\sigma^2})$, and changes the noise, $< \epsilon^2 >^{\frac{1}{2}} \rightarrow < \epsilon^2 >^{\frac{1}{2}} \div \sqrt{2\pi\sigma_g^2}$. A similar

† It is straightforward to generalise to non-independent noise using a generic multivariate Gaussian distribution, which leads to a "least-squares" estimator based on a Mahanoublis metric.

‡ This is the same operation as convolution for symmetric profiles – desirable since they induce no systematic shift in image position. Most detection filters are also normalised to unit area under the curve, to ensure average intensities remain the same.

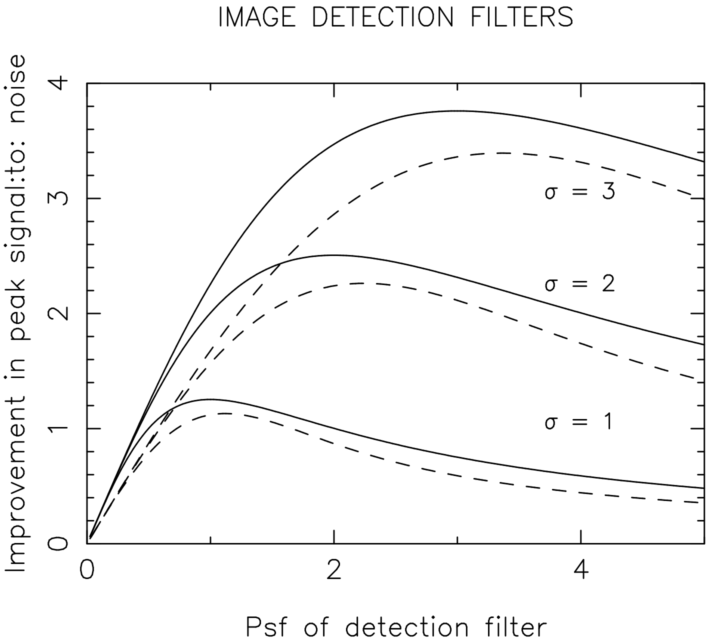

FIGURE 2. Performance of various detection filters for target Gaussian images of $\sigma = 1, 2, 3$. Solid lines are for Gaussian filters and dashed lines are for circular top-hat filters. Note: how the improvement peaks for filters of equivalent scale size to the target objects; how insensitive the improvement is to the exact size of the filter; and how insensitive the performance is to the exact shape of the filter.

---

exercise can be carried through for normalised circular top-hat filters of radius, a, to show that the improvement becomes $\sqrt{\pi a^2} \times \sigma^2/a^2 \times exp(-a^2/\sigma^2)$. Because the performance gain in peak signal:to:*rms* noise is relatively insensitive to the exact size and shape of the filter, it is then feasible to use non-linear variants of the preceding filters to detect large LSBGs in the presence of contaminating foreground/background objects. See Irwin *et al.* (1990) for more details and examples of the galaxies detected.

2.3. *Detecting Using Pixel Connectivity*

The main purpose of the detection filter is to enhance the peak signal:to:noise to optimise the performance of object locating algorithms. Most of these algorithms are some variant on the theme of peak (eg. Herzog & Illingworth, 1977) or island searches (eg. FOCAS – Tyson & Jarvis, 1979). Both methods work by first prescribing a pixel detection threshold above background, otherwise known as a detection isophote, and then proceed to look for either isolated peaks, or islands of contiguous pixels above this threshold. Since most peak-based methods examine the intensities of neighbouring pixels too, there is in practice, little difference between the two approaches and I will stick to describing the pixel connectivity approach.

Figure 3 illustrates the main concepts. Potential images are defined as regions of simply-connected pixels above the detection isophote, relative to the local sky back-

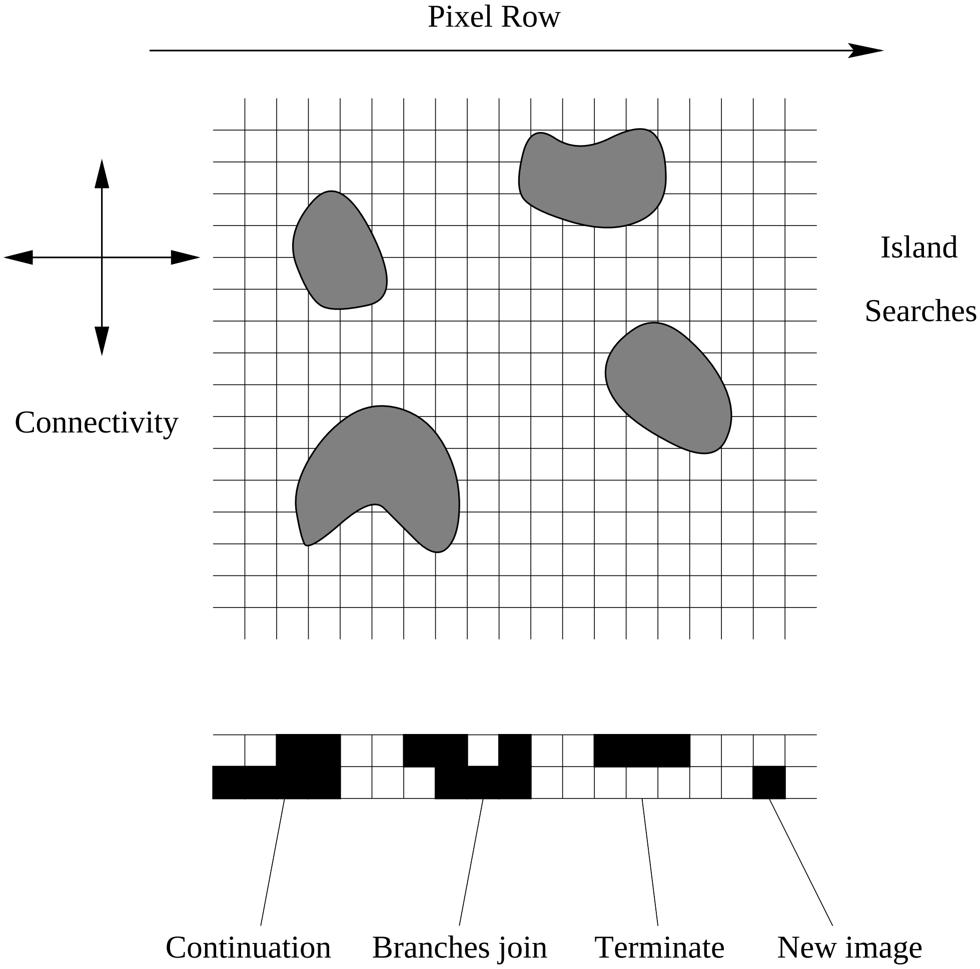

Figure 3. The basis of pixel connectivity algorithms

ground. Two passes through the data are recommended. On the first pass the sky background is estimated, as outlined previously, and on the second pass the sky level (interpolated if necessary) plus the threshold, define those pixels that require further examination. Spurious images or noise are rejected using a combination of the requirements for a pixel to be above threshold and for it to be connected with enough neighbours to meet a chosen minimum size criterion. The threshold is usually set to be some fixed multiple ($\approx$2) of sky noise above sky. Once a potential image has been located, the original pre-detection filter intensities of the pixels are used for subsequent analysis. This latter requirement ensures that any derived image parameters are unaffected by the blurring action of the detection filter. Of the two most commonly used isophotal connectivity paths, the four nearest neighbour route (N,S,E,W) seems to give better performance than the eight-fold neighbour approach (N,S,E,W + NW,SE,NE,SW). The latter type of connectivity tends to produce more spurious pixels in the outer parts of an image leading to somewhat noisier image parameters.

The operations required to do all this, using only one pass through the data and minimal memory requirements, ie. enough rows of pixels in memory to do the detection filtering, can be summarised as follows (see Lutz, 1979 for more details):

(*a*) Set up parent, link-list and pixel information arrays/stack. The parent list, a label for potential new images, points to the position of the 1st link in the link-list for this image; the link-list entry then points to the next pixel for this image and so on; while the pixel information arrays store the i,j position and original intensity of the pixel in the location specified within the link-list.

(*b*) Read in a new row of pixels; do filtering and save original intensities. The filtering is of necessity out of phase with the row of data currently being analysed for objects.

(*c*) Set the background level and add on threshold and flag pixels above this level.

(*d*) Is pixel i,j neighbour to any active pixels in the current parent/link lists or not ? It must be either: a continuation of an image to add to the linked-list; a joining of two branches of previously separate images (see figure 3) requiring updates of parent and link-list arrays; or a potential new image.

(*e*) If within a row no pixels join to a previously active parent image, terminate that image, process pixels to generate parameters and return free slots to stack.

2.4. *Estimating Image Parameters*

The isophotal connectivity method of the previous section only makes use of the information stored in those pixels that passed the threshold and contiguity criteria. The advantages of this approach are that it is computationally fast and straightforward to implement; and that a wide variety of useful and robust image parameters can be derived. The following list demonstrates the range of typical image parameters that can be computed from the connected pixel list in a single pass. These are the image parameters that are widely used by the measuring machine community for astrometry, photometry and image shape description (see for example MacGillivray & Stobie, 1985). It is worth stressing that a single photographic Schmidt plate, the raw material for the measuring machines, contains up to 4 Gbytes of data, recorded in a highly non-linear fashion, and that the currently available sky surveys total well over 20 Tera-bytes of data. The ability to process large volumes of data quickly is also highly relevant for processing wide field CCD data. Examples of alternative strategies, usually involving significantly more computational effort, will be given in the next section. In all the following equations, $I(x_i, y_i)$ denotes the intensity of the ith pixel found in the island search phase and the sum is over all detected connected pixels. We shall assume for simplicity that all detected images are single objects. For more details on dealing with overlapped blended images see Irwin (1985).

Isophotal Intensity – the integrated flux within the boundary defined by the threshold level; ie. 0th moment

$$I_{iso} = \sum_i I(x_i, y_i) \tag{2.21}$$

For Gaussian images, this is related to the total intensity by the factor $(1 - I_t/I_p)^{-1}$, where I_p is the peak flux and I_t the threshold level relative to sky. This implies that for images close to the detection limit, $I_p \approx I_t$, isophotal magnitudes will be around 1 magnitude fainter than total magnitudes.

Position – computed as an intensity-weighted centre of gravity; ie. 1st moment

$$x_o = \sum_i x_i.I(x_i, y_i) / \sum_i I(x_i, y_i) \tag{2.22}$$
$$y_o = \sum_i y_i.I(x_i, y_i) / \sum_i I(x_i, y_i)$$

Note that the intensity weighting and the isophotal connectivity criteria automatically reduce the sensitivity of the estimator to outlying points.

Covariance Matrix – the triad of intensity-weighted 2nd moments is used to estimate the eccentricity/ellipticity, position angle and intensity-weighted size of an image

$$\sigma_{xx} = \sum_i (x_i - x_o)^2.I(x_i, y_i) / \sum_i I(x_i, y_i) \tag{2.23}$$
$$\sigma_{xy} = \sum_i (x_i - x_o).(y_i - y_o).I(x_i, y_i) / \sum_i I(x_i, y_i)$$
$$\sigma_{yy} = \sum_i (y_i - y_o)^2.I(x_i, y_i) / \sum_i I(x_i, y_i)$$

The simplest way to derive the ellipse parameters from the 2nd moments is to equate them to an elliptical Gaussian function having the same 2nd moments. It is then straighforward to show (eg. Stobie, 1980) that the scale size, $\sqrt{\sigma_{rr}}$, is given by, $\sigma_{rr} = \sigma_{xx} + \sigma_{yy}$; the eccentricity, $ecc = \sqrt{(\sigma_{xx} - \sigma_{yy})^2 + 4.\sigma_{xy}^2} \ / \ \sigma_{rr}$; and the position angle, θ is defined by, $tan(2\theta) = 2.\sigma_{xy}/(\sigma_{yy} - \sigma_{xx})$. Higher order moments have been used as shape descriptors (eg. Tyson & Jarvis, 1979), but are generally much more prone to corruption by outlying noisy data pixels.

Areal Profile – an interesting variant on the radial profile, which measures the area of an image at various intensity levels. Unlike the radial profile, which needs a prior estimate of the image centre, the areal profile provides a single pass estimate of the profile

$$Areal\ Profile \ \rightarrow \quad T + p_1, T + p_2, T + p_3, ...T + p_m \tag{2.24}$$

where $p_j; j = 1, ...m$ are intensity levels relative to the threshold, T, usually spaced logarithmically to give even sampling.

The peak height, I_p, is a useful related addition to the areal profile information and is defined as

$$I_p = max[I(x_i, y_i)]_i \tag{2.25}$$

or alternatively measured by extrapolation from the areal profile if the image is saturated (straightforward for photographic plates, but charge leakage causes problems with CCD data).

Other useful measures of image size are those defined by Kron (1980) based on moments of the radial profile.

$$r_1 = \int r.I(r)dr / \int I(r)dr; \quad r_{-2} = \int r^{-2} I(r)dr / \int I(r)dr \tag{2.26}$$

These involve more than one pass through the data but are good shape discriminators for morphological classification. There are many other methods of parameterising the shape information as the front end to an object classification scheme (eg. star/galaxies/noise). A recent review of the field is given by Odewahn (1995).

3. Alternatives Methods for Astrometry and Photometry

Maximum likelihood theory enables us to devise optimum methods for estimating image parameters and also specifies the minimum attainable error it is possible to achieve for images of a given profile and signal:to:noise (eg. Irwin, 1985). Isophotal methods are not generally optimal, but they are robust, require no prior knowledge of the image profile, and are simple to compute. A serious drawback in their direct use is that the derived intensities are biased, since a varying fraction of the total flux is included. Although there are assorted tricks to get around this problem for stellar images (eg. Bunclark & Irwin, 1983), recovering the lost flux for an unknown galaxy profile is more difficult. From an information-theoretic viewpoint the loss of information in neglecting the outer parts of the profile is not so serious, since most of the information regarding image location and total intensity comes from the higher signal:to:noise inner part of the profile. Consequently, bias notwithstanding, isophotal methods actually come quite close, certainly within <50%, of the minimum attainable error bound.

In principle, a stellar image is completely specified by its coordinates and intensity. In practice, variations in psf across the field and lack of prior knowledge about which images are stellar conspire to make life more complicated. The problem is much worse for galaxies, since we have no prior knowledge of their profile. These problems have lead to a whole range of methods, particularly photometric methods, to circumvent the lack of knowledge.

3.1. *Aperture Photometry*

There are two basic approaches for estimating total intensities. The first is to integrate the flux within an aperture. The second is to fit some sort of profile to the image and then renormalise the profile-based intensity to be equivalent to a total one. Apertures are either fixed in size (usually for bright stellar photometry only) or allowed to "grow" automatically out to some pre-specified maximum area.

Aperture photometry is defined by

$$I_{ap}(r) = \sum_{i}^{N} I_i - N \times sky \tag{3.27}$$

such that all pixels indexed by i are within a radius, r, of the image centre †. An isophotal estimate of the image centre, as in equation 2.22, is usually adequate. Intensity contributions from pixels crossing the boundary are usually added in on a pro-rata basis. Defining I_i to be the object flux + sky flux in pixel i, clarifies the dependence of $I_{ap}(r)$ on the accuracy of the sky estimate – assumed constant over the image. The problem with aperture photometry is the necessary tradeoff between capturing all the light, = large radius; adding in too much sky noise = $\sqrt{N} \times \sigma_{sky}$; and being mislead by systematic errors in the sky level = $N \times \Delta_{sky}$. The usual approach to minimise these problems is to "examine" the curve-of-growth, $I_{ap}(r)$ -v- r, as illustrated in figure 4. Another problem, not shown in the figure, is the presence of neighbouring images, which will also lead to an upturn in the curve. In this example, the systematic errors in the sky estimate correspond to ±1% errors in the absolute sky level for a typical scale length galaxy image at a similar magnitude to a 1 arcsec^2 patch of sky. Figure 2. of Irwin (1985) shows how typical estimators of sky can readily be off by 1% in even moderately crowded fields.

For stellar images, assuming the psf does not vary rapidly around the field, it is possible

† Elliptical apertures, parameterised by semi-major axis radius, were used by Irwin & Hall (1982) for galaxy photometry to minimise the area of sky included.

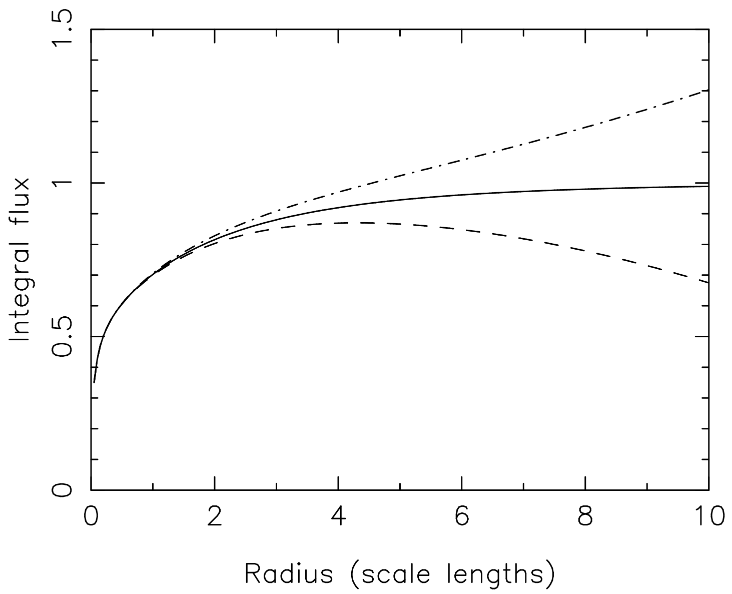

FIGURE 4. Idealised curves-of-growth for aperture integration of an exponential profile. If the sky level is exactly correct, $I_{ap}(r)$ will asymptotically stabilise to the desired level. If the sky estimate is too high the curve will turn over (dashed), too low and the curve continues rising (dot-dash). The effect of noise in the data is not shown.

to minimise the preceding problems by calibrating the curve-of-growth using isolated bright stellar images. Fainter stars can then be integrated out to a compromise radius (usually a few core radii) and then a well defined correction applied. Clearly this is not possible for galaxy images and some sort of empirical fit to the curve-of-growth (eg. 3rd order polynomial and look for zero derivative out to some maximum radius) seems to be as good a method as any of automating the decision and minimising the effects of sky noise (see Irwin & Hall (1982) for more details).

3.2. *Profile Fitting*

Profile fitting, usually limited to stellar images, is now well served by several software packages such as: ROMAFOT – Buonanno & Iannicola (1989); DoPHOT – Mateo, Saha & Schecter (1990); and DAOPHOT – Stetson (1987). These references address the main problems and issues involved in successful profile fitting; whilst the work by King (1983), Irwin (1985) and Goad (1986) investigates the theory underpinning optimal fitting and shows how this can be used to predict the likely errors of the derived parameters. Stellar psf fitting is the only reliable technique to use in crowded regions such as those found toward the centre of globular clusters. By iterative use of difference maps (data – psf model) and by simultaneous fitting of overlapping components in a blended image, results very close to the theoretical error bound can be achieved.

Very briefly, the main decisions to be made in profile fitting are:

(a) is the profile, ie. psf, adequately constant across the frame or does some allowance for geometric variation need to be folded in (eg. HST data; ROSAT data) ?

(b) can the psf be approximated using an analytic function ? Usually some mixture of Gaussian-exponential-Lorentzian or a Moffat profile, $(1 + r^2/a^2)^{-\beta}$, are used.

(c) if the images are noticeably trailed is an elliptical deformation of one of the above profiles adequate ? or,

(d) does an empirically constructed profile need to be formed from isolated stars ? If the seeing psf is not well sampled then some sort of sub-pixel interpolation will be necessary to reconstruct the model profile at non-integral pixel locations. This is easy with analytic profiles, harder for empirical profiles.

(e) how large a region around each image should the psf be fitted over ? With a complex blend of several images this is not necessarily straightforward to answer. A simple solution is to only use those pixels picked out by some isophotal detection algorithm.

(f) should a local estimate/update for the background level, caused by unresolved images etc., be made simultaneously with the psf parameter fitting ?

(g) if the profile is an exact match for the psf, total intensities = profile intensities. In practice this is rarely the case and isolated bright stars are used to estimate the offset between the profile and total intensity systems.

For more details on these and other points consult the preceding references.

There is a very simple link between isophotal intensity estimation and profile fitting for isolated images, which goes some way to show why isophotal estimates are so popular. With the usual Gaussian/Poisson noise encountered in optical work, all profile fitting methods reduce to weighted least-squares methods of the form of minimising

$$F = \sum_i \frac{(d_i - m_i)^2}{\sigma_i^2} \tag{3.28}$$

where d_i and m_i are respectively the data and model at pixel i, and σ_i^2 is the expected noise variance $< \epsilon_i^2 >_i$, such that $d_i = m_i + \epsilon_i$. † For a single profile, with assumed known position, $m_i = k.\phi_i$, where k is the unknown scaling factor and ϕ is the normalised psf (ie. $\int\int \phi(x,y)\ dxdy = 1$) shifted to the appropriate location. Solving for the intensity parameter then gives

$$k = \frac{\sum_i \phi_i.d_i/\sigma_i^2}{\sum_i \phi_i^2/\sigma_i^2} \tag{3.29}$$

which is nothing more than a weighted sum over the pixel intensities belonging to the image; although subtle differences between the data–model overlap integral in the numerator, and the model–model overlap integral in the denominator, are the cause of the intensity renormalisation problem mentioned previously. There are two interesting limiting cases: for bright images well above sky, Poisson noise dominates, so that $\sigma_i^2 = \phi_i$ (if the gain is unity, $\propto$ otherwise) and the profile intensity is simply $\sum_i d_i$, the isophotal intensity; for a faint image dominated by constant sky noise the profile intensity is just the profile-weighted sum of the pixel intensities. This analogy can be extended to multiple overlapping images, using more complex weighting factors.

3.3. *Astrometry – Image Centering*

In the late 1970's the most popular method for locating the centre of an image involved using marginal sums of the image intensity distribution. This reduces a two-dimensional

† For those condemned to work with few photons per resolution element, as is often the case in X-ray or γ-ray astronomy, the equivalent expression is, $F = \sum_i d_i.ln(m_i) - m_i - ln(d_i!)$.

problem to a pair of one-dimensional problems, with subsequent savings in computer time. The marginal sum method involves some variant along the lines of the following steps. First, form $\rho(x_i)$, and equivalently $\rho(y_i)$, where

$$\rho(x_i) = \sum_{j=-a}^{j=+a} I_{i,j} \tag{3.30}$$

and then find the centre of the one-dimensional distribution(s) using either: the median, a Gaussian fit to the distribution (often including a sloping background component rather than a constant), or a variant of the centre of gravity such as

$$\bar{x} = \frac{\sum_i [\rho(x_i) - \bar{\rho}].x_i}{\sum_i [\rho(x_i) - \bar{\rho}]} \tag{3.31}$$

using only those pixels with $\rho(x_i) \geq \bar{\rho}$. Iterate through these two steps if $\bar{x}$ is more than 1 pixel from the window centre. Although computationally fast, there are several shortcomings of the method: the window has to be centred fairly accurately to start with or the iterative solution can diverge; neighbouring images can corrupt the marginal distribution and bias the answer away from the true centre; if the window $[-a, a]$ is too large excessive noise is folded in, too small and much of the information is lost. Irwin (1985) outlined how this potential loss of information degrades the performance of the algorithm and suggested ways of minimising the problem by using Maximum Likelihood theory to establish the optimum window to use.

In many ways, the marginal sum method was superceded by the centre-of-gravity estimates derived from pixel connectivity algorithms – ie. isophotal centre-of-gravity estimators such as equation 2.22. If we generalise such an estimator to include a weighting term, w_i, (cf. equation 3.29 for intensity estimation) along the lines of

$$\bar{x} = \frac{\sum_i w_i \times x_i \times I_i}{\sum_i w_i \times I_i} \tag{3.32}$$

where the summation is over all isophotally connected pixels, is there an optimum weighting factor ? Put another way, what does Maximum Likelihood theory suggest ? For an image with underlying normalised profile, $\phi(x, y)$, the optimum weight is given by

$$w_i = \left[\frac{1}{x - \bar{x}} \times \frac{\delta\phi}{\delta x}\right]_i \times \frac{1}{\sigma_i^2} \tag{3.33}$$

Clearly, this involves an iterative solution for $\bar{x}$ (and $\bar{y}$). Irwin (1985) shows how to extend this scheme to overlapping blended images.

As an example, for a radially symmetric Gaussian profile, this simplifies to $w_i \propto \phi(x_i, y_i)/\sigma_i^2$. As in the previous section, the two limiting cases are very revealing. For bright images, w_i is a constant, implying that for such cases isophotal centre-of-gravity methods are close to optimal. For faint images, $w_i \propto \phi(x_i, y_i)$, and profile weighting, with iteration, is the answer.

For more details of the derivation of optimum estimators and a full error analysis see the papers by, King (1983), Irwin (1985) and Goad (1986).

4. Acknowledgements

In preparing the section on telescope optics I have benefited from many fruitful conversations with Roderick Willstrop and Ian Parry and would like to thank them for their help. I am also indebted to Richard McMahon for letting me borrow from his notes of sky surveys in other wavebands. Finally I would like to acknowledge the financial support provided by the Instituto de Astrofisica de Canarias, and to thank them for inviting me to take part in the Winter School.

REFERENCES

Becker, R.H., White, R.L. & Helfand, D.J., 1995, *Astrophys. J.,* **450**, 559.

Bijaoui, A., 1980, *Astr. Astrophys.,* **84**, 81.

Boksenburg, A., 1972, in *Proceedings of the ESO-CERN Conference on Auxiliary Instrumentation for Large Telescopes*, eds. S. Lausen & A. Reiz, 295.

Boronson, T., Reed, R., & Wong, W.Y., 1994, in *Instrumentation in Astronomy VIII*, Proc. SPIE, **2198**, 877.

Bradt, H.V.D., Ohashi, T., & Pounds, K.A., 1992, in *Ann. Rev. Astron. Astrophys.*, **30**, 391.

Buonanno, R., & Iannicola, G., 1989, *Publ. astr. Soc. Pacific,* **101**, 294.

Bunclark, P.S., & Irwin, M.J., 1983, in *Proc. Statistical Methods in Astronomy*, ESA SP-201, 195.

CFHT MOCAM User Manual, 1995.

Dainty, J.C. & Shaw, R., 1974, in *Image Science*, Academic Press.

Demers, S., Kibblewhite, E.J., Irwin, M.J., Nithakorn, D.S., Beland, S., Fontaine, G., & Wesemael, F., 1986, *Astr. J.,* **92**, 878.

Fisher, R.A., Corbett, A.S., & Williams, C.B., 1943, Journal of Animal Ecology, **12**, 42.

Fordham, J.L.A., Bone, D.A., Michel-Murillo, R., Norton, T.J., Butler, I.G., & Airey, R.W., 1994, in *Instrumentation in Astronomy VIII*, Proc. SPIE, **2198**, 829.

Geary, J.C., Luppino, G.A., Bredthauer, R., Hlivak, R.J. & Robinson, L., 1991, *Proceedings of The SPIE: CCDs and Solid State Optical Sensors II*, **1447**, 264.

Gilliland, R.L., 1992, in *Astronomical CCD Observing and Reduction Techniques*, ed. S.B.Howell, ASP Conf. Series, Vol. 23, 68.

Gilmore, G., & Walton, N., 1993, *Gemini*, **40**, 16.

Goad, L., 1986, in *Instrumentation in Astronomy*, SPIE Vol. **627**, 688.

Gullixson, C.A., 1992, in *Astronomical CCD Observing and Reduction Techniques*, ed. S.B.Howell, ASP Conf. Series, Vol. 23, 130.

Gunn, J.E., *et al.* , 1987, *Optical Engineering*, **26**, 779.

Hanisch, R.J., 1992, in *Astronomical CCD Observing and Reduction Techniques*, ed. S.B.Howell, ASP Conf. Series, Vol. 23, 285.

Hartley, M., 1994, in *Digitised Optical Sky Surveys* , eds. H.T.MacGillivray and E.B. Thomson, Kluwer Academic Publishers, Dordrecht, p.117.

Herzog, A.D., & Illingworth, G., 1977, *Astrophys. J. Suppl.,* **33**, 55.

Hoaglin, D.C., Mosteller, F., & Tukey, J.W., 1983, *Understanding Robust and Exploratory Data Analysis*, Wiley, Chp 11.

Irwin, M.J., & Hall, P., 1982, *Occ. Repts. R. Obs. Edin.*, **10**, 111.

Irwin, M.J., & Trimble, V., 1984, *Astr. J.,* **89**, 83.

Irwin, M.J., 1985, *Mon. Not. R. astr. Soc.,* **214**, 575.

Irwin, M.J., Davies, J.I., Disney, M.J., & Phillipps, S., 1990, *Mon. Not. R. astr. Soc.,* **245**, 289.

Irwin, M.J., 1992, in *Digitised Optical Sky Surveys* , eds. H.T.MacGillivray and E.B. Thomson, Kluwer Academic Publishers, Dordrecht, p.43.

Irwin, M.J., 1994, in *Working Group on Wide Field Imaging – Newsletter No. 5*, 25.

Janesick, J., & Elliot, T., 1992, in *Astronomical CCD Observing and Reduction Techniques*, ed. S.B.Howell, ASP Conf. Series, Vol. 23, 1.

Jorden, P.R., Deltorn, J., & Oates, A.P., 1994 in *Instrumentation in Astronomy VIII*, Proc. SPIE, **2198**, 836.

Kibblewhite, E.J., Bridgeland, M.T., Bunclark, P.S., & Irwin, M.J., 1984, in *Proc. Astronomical Microdensitometry Conference*, NASA-2317, 277.

King, I., 1983, *Publ. astr. Soc. Pacific,* **95**, 163.

Kron, R.G., 1980, *Astrophys. J. Suppl.,* **43**, 1.

Kron, R.G., 1995, *Publ. astr. Soc. Pacific,* **107**, 766.

Lasker, B.M., 1995, in *Future Utilisation of Schmidt Telescopes*, IAU Symposium No. 148, eds. R.D.Cannon *et al.* , ASP, San Francisco (in press).

Lasker, B.M., 1995b, *Publ. astr. Soc. Pacific,* **107**, 762.

Lesser, M.P., 1989, in *CCDs in Astronomy*, ed. G.H. Jacoby, ASP Conf. Series, Vol. 8, 65.

Lesser, M.P., 1994, in *Instrumentation in Astronomy VIII*, Proc. SPIE, **2198**, xxx.

Luppino, G.A., Bredthauer, R.A., & Geary, J.C., 1994 in *Instrumentation in Astronomy VIII*, Proc. SPIE, **2198**, 810.

Lutz, R.K., 1979, in *Image Processing in Astronomy*, Trieste, eds. G. Sedmak, M. Capaccioli & R. J. Allen, 218.

MacGillivray, H.T., & Stobie, R.S., 1985, in *Vistas in Astronomy*, **27**, 433.

Mateo, M., Saha, P., & Schecter, P., 1990, DoPHOT Users Manual.

Morgan, D.H., 1995, in *Future Utilisation of Schmidt Telescopes*, IAU Symposium No. 148, eds. R.D.Cannon *et al.* , ASP, San Francisco (in press).

Newell, E.B., 1982, *Occ. Repts. R. Obs. Edin.*, **10**, 111.

Odewahn, S.C., 1995, *Publ. astr. Soc. Pacific,* **107**, 770.

Parker, Q.A., Philipps, S., Morgan, D.H., Malin, D.F., Russell, K.S., Hartley, & Savage, A., 1994, in *Astronomy from Wide-Field Imaging*, IAU Symposium No. 161, eds. H.T.MacGillivray *et al.* , Kluwer Academic Publishers, Dordrecht, p.129.

Pratt, W.K., 1978, *Digital Image Processing*, Wiley.

Prestage, R., & Peacock, J., 1983, *Mon. Not. R. astr. Soc.,* **204**, 355.

Price, S., 1988, *Publ. astr. Soc. Pacific,* **100**, 171.

Russell, K.S., Malin, D.F., Savage, A., Hartley, M. & Parker, Q.A., 1992, in *Digitised Optical Sky Surveys* , eds. H.T.MacGillivray and E.B. Thomson, Kluwer Academic Publishers, Dordrecht, p.23.

Seidel, L., 1856, Astronomische Nachrichten, **43**, Nos. 1027,1028,1029.

Sekiguchi, M., Iwashita, H., Doi, M., Kashikawa, N., & Okamura, S., 1992, *Publ. astr. Soc. Pacific,* **104**, 744.

Smart, W.M., 1956, *Spherical Astronomy*, CUP, Cambridge.

Stetson, P.B., 1987, *Publ. astr. Soc. Pacific,* **99**, 191.

Stobie, R.S., 1980, *Journal British Interplanetary Society*, **33**, 323.

Stubbs, C.W., *et al.* , 1994, in *Charge Coupled Devices and Solid State Optical Sensors III*, Proc. SPIE, 1900.

Thejll, P., & Shipman, H.L., 1988, *Publ. astr. Soc. Pacific,* , **100**, 398.

Tulloch, S., 1995, *ING La Palma Technical Note No. 99.*

Tukey, J.W., 1971, in *Exploratory Data Analysis*, Adison Wesley.

Tyson, J.A., & Jarvis, J.F., 1979, *Astrophys. J.,* **230**, L153.

Tyson, J.A., Bernstein, G.M., Blouke, M.M., & Lee, R.W., 1993, in *High-Resolution Sensors and Hybrid Systems*, Proc. SPIE, 1656, 400.

Willstrop, R.V., 1985, *Mon. Not. R. astr. Soc.,* **216**, 411.

Woodward,.....& Davies,.... 1958............

York, D., *et al.* , 1993, A Digital Sky Survey of the Northern Galactic Cap, NSF proposal.