Centroiding Images in the Photometric Pipeline

Robert Lupton and Jim Gunn

August 26, 1998

1 Introduction

This document discusses methods for finding the centres of objects, with main emphasis upon stars. Both saturated and non-saturated objects are considered.

2 Astrometric Centering: Gaussian Quartic Interpolation Schemes

It was apparent in the earlier document that simple quadratic interpolation was pretty good, but not quite good enough, incurring systematic errors of the order of 10-15 mas depending on the location of an image with respect to the pixel grid. While random errors of this sort are barely acceptable for the measurements of individual objects (we expect at best to have 30 mas errors from seeing) they are not acceptable for either the secondary or primary astrometric standards, which define the coordinate system. A simple *ad hoc* correction was described there that had no basis in theory but which improved the error by a factor of two and a half or so. A much better one has been found recently in the process of analyzing the situation properly, and that analysis is presented here.

2.1 The Interpolator

The key is very simple; one knows that the image near the peak is roughly gaussian; the ideal smoother is a gaussian of approximately the same sigma. Our standard seeing model for the SDSS is a two-component gaussian of sigmas 0.95 and 2.7 pixels and amplitude ratio 9:1. This has a FWHM of 1.0 arcsec, and though the flux is roughly equally shared by the two components, essentially all the high-frequency power which is useful for astrometric centering is contained in the small one. We simply make use of the near-gaussianity of the image to find the 4th-order correction term to the simple quadratic interpolator (which, of course, itself emerges from expanding the gaussian to second order only.)

Let us look at the problem in one dimension only at the beginning. Consider a triad of points on which the function to be interpolated has values v_-,v_0, and v_+; we know that v_0 is at least as large as the other two. The abscissae x_-, x_0, and x_+ are -1, 0, and +1, and we wish to find the value of x at the maximum, supposing that the function is a gaussian. Let the abscissa at the maximum be simply x.

Then

$$v_0 = A \ \exp\left(\frac{-x^2}{4\sigma^2}\right) = A\left(1 - \frac{x^2}{4\sigma^2} + \frac{x^4}{32\sigma^4} + \ldots\right) \tag{1a}$$

$$v_+ = A \ \exp\left(\frac{-(1-x)^2}{4\sigma^2}\right) = A\left(1 - \frac{(1-x)^2}{4\sigma^2} + \frac{(1-x)^4}{32\sigma^4} + \ldots\right) \tag{1b}$$

$$v_- = A \ \exp\left(\frac{-(1+x)^2}{4\sigma^2}\right) = A\left(1 - \frac{(1+x)^2}{4\sigma^2} + \frac{(1+x)^4}{32\sigma^4} + \ldots\right), \tag{1c}$$

where we have parametrized the gaussian, which in principle is arbitrary, with the form expected for a star image smoothed with a similar gaussian — i.e., σ in the above expressions is the σ associated with the central gaussian in the star image. A is the amplitude in the *smoothed* image, which is related to the amplitude a in the unsmoothed image by $A = a/2$; this would be the case, at least, if the stars were single-component gaussians of width σ. In reality, to the extent that the 2-component model is correct, the amplitude of the small component is reduced by a factor of 2 and that of the large component hardly at all (about 5 percent). Thus the amplitude of the small component in the smoothed image is roughly

$$A \approx 0.45a,$$

and this is superimposed upon a gaussian of a bit more than twice this width (the ratio is $\sqrt{2}/\sqrt{10}$) with amplitude $\sim .095a$.

Simple algebra using these relations produce the following:

$$\frac{\sigma^2(v_+ - v_-)}{A} = x - \frac{x(1+x^2)}{4\sigma^2} \tag{2}$$

and

$$\frac{\sigma^2(2*v_0 - v_+ - v_-)}{A} = \frac{1}{2} - \frac{1+6x^2}{16\sigma^2} \tag{3}$$

The second terms on the RHS of each of these expressions come from the fourth-order terms in the expansions—it is clear that the expansion parameter here is x/σ, and that the expansions are very good for large σ; the range of x is clearly ± 0.5. If we drop those terms, we retrieve the standard quadratic interpolator,

$$x = \frac{v_+ - v_-}{2(2v_0 - v_+ - v_-)} = \frac{s}{d_2}, \tag{4}$$

where $s = (v_+ - v_-)/2$ is the mean slope across the array, and $d_2 = (2v_0 - v_+ - v_-)$ is the second difference.

We also get an estimate for σ^2;

$$\sigma^2 = \frac{A}{2d_2}. \tag{5}$$

If we use the fourth-order terms to form an improved estimate for σ^2, and use the second-order estimates whenever appropriate in the 4th-order correction terms, we obtain

$$\sigma^2 = \frac{A}{2d_2} - \frac{(1+6*x^2)}{8}, \tag{6}$$

and

$$x = \frac{s}{d_2} + \frac{x(1-4x^2)}{8\sigma^2}. \tag{7}$$

Note that the correction term has the correct 'symmetry' properties; it vanishes at $x = 0$ and $x = 1/2$, as it must. We can write it simply in terms of differences (again using the second-order estimates for quantities in the correction term):

$$x = \frac{s}{d_2}\left(1 + \frac{s_+ s_-}{Ad_2}\right) \tag{8}$$

where $s_+ = v_0 - v_+$, and $s_- = v_0 - v_-$. Now, in reality, the PSF is not exactly gaussian, and for the parameters of interest ($\sigma \approx 1$) the fourth-order and higher terms are not *so* small, so the amplitude of the correction term for best results may not be exactly what this simple development would predict. In fact, we find later that one needs a term about a third larger to produce the best astrometric results; i. e.

$$x = \frac{s}{d_2} \left(\frac{1 + k s_+ s_-)}{A d_2} \right) \tag{9}$$

with $k \approx 1.33$. Upon investigation, it becomes clear that this term arises more-or-less directly from the fact that the second-order estimate for σ^2 (Equation 5) is too large by about this factor. Most of this, in turn, comes from two sources: The fourth-order correction to σ^2 (Equation 6) is a negative term of order 15 percent (all these remarks are for $\sigma \approx 1$), and the amplitude A one measures is contaminated by the much flatter large gaussian for our assumed two-component PSF, so its effective value is about 20 percent smaller than the whole amplitude. It is important to note that the errors if one ignores all of this and uses the calculated 4th-order corrections as is are of order 3-4 mas in one dimension, which is quite acceptable. On the other hand, one may as well use the more accurate empirical form.

The correction term has a maximum at $\sqrt{(1/12)} = 0.288$, and the size there (with $\sigma \approx 1$) is $.024k$. With our 400 mas pixels, the maximum correction is 12 mas, certainly of the right order to 'fix' the quadratic error.

These formulae (and the similarly constructed ones for y) will serve for the astrometric standards and bridge stars; for the photometric pipeline, there are statistical considerations for very faint objects which we will consider anon.

2.2 Extension to Two Dimensions; Accuracy Attainable With the SDSS

The extension of this interpolator to two-dimensional images is not completely straightforward. In those cases in which the images are round or elongated (by aberrations, tracking errors, or whatever) in such a manner that the major axes of the elongation are along the coordinate axes, a simple extension of the technique using either the x and y stripes through the maximum pixel in the smoothed image or suitably interpolated stripes to better 'center' the data on the supposed real center works satisfactorily. In general, however, this condition is not met for imperfect images, and they are elongated at some arbitrary angle to the axes. The *ad hoc* crossstripe methods fail quite badly for these images, incurring errors of up to 15 mas or so; the errors can easily exceed the quadratic ones, even for images which are quite accurately gaussian near the peak.

It is the case, however, that a two-dimensional gaussian of arbitrary ellipticity and inclination has the property that any one-dimensional cut through it is also a gaussian, which leads immediately to a two- dimensional extension which works very well indeed. Consider point which is the maximum in the smoothed image and its eight neighbors. Call the values of the smoothed image at these points $V_{i,j}, -1 \leq i,j \leq 1$; The central point has value $V_{0,0}$. Then the three triads in x for the three rows can each be interpolated (since they are approximately three values of some gaussian) using the quartic interpolator, to find three maxima mx_-, mx_0, and mx_+ along the lines $y = -1, 0, +1$. Likewise the three triads along the columns for the three x values can be interpolated to find the three maxima my_- ,my_0, and my_+ belonging to $x = -1, 0, +1$. The mx's lie on a curve along which the two-dimensional maximum must lie, as do the my's, and indeed the maximum must lie at the intersection of those curves. In real application to simulated SDSS images, the sets of three maxima lie essentially on straight lines (which is strictly the case for a real two-dimensional gaussian), and one could either use the least-squares straight line and find the intersection analytically, or, as we have done, allow the three points to define a quadratic and do a simple iteration to find the intersection: Beginning with the quadratic position for x (y), solve for the y (x) maximum; using those estimates for x (y), solve again for y (x). The slopes are quite small, and one iteration suffices. The accuracy of this technique is excellent; the real gaussian least-squares center is returned to within 1.7 mas for every location on every chip in the photometric and astrometric array using the full simulated TDI images for the as-designed optical system.

It is worth emphasizing that the limitations of the gaussian fitting itself are a bit worse than this. The peak errors amount to about 5 mas in the worst case—i.e., the difference between the least-squares gaussian fitted center and the real centroid of the image is no larger than about 5 mas anywhere in the array, is considerably smaller in the chips away from the periphery, but rises to about this value somewhere in most of the edge detectors. These errors, however, vary smoothly over the field and are unaffected, of course, by where in a *pixel* a star lands; they are also much smaller than the residual geometric distortions of the design, up to about 80 mas over some of the chips in the array. Some of this is removed by simple scale changes, which will be solved for anyway, but about 30 mas remains, compared to about 2 mas for the gaussian fitting error. Both terms are essentially perfectly removed (to less than 1 mas) by cubic corrections separately in x and y (in TDI data there cannot be crossterms.) The optical contribution can probably be adequately dealt with using the ray-trace calculations; it is possible that the calculations used here could adequately predict the gaussian errors as well, though at the 2 mas RMS level it is not clear that one wishes to bother.

2.3 Statistics

Let us assume that the original picture has pixel values $p_{\mathbf{i}}$; the subscript $\mathbf{i}$ is a two-dimensional object (x,y), of course. We assume that the noise in the image is such that the noise in distinct pixels is independent. The smoothed image in which one searches for maxima to find positions is generated by

$$P_{\mathbf{i}} = \sum w_{\mathbf{i-j}} * p_{\mathbf{j}}. \tag{10}$$

We assume that the weights are normalized so that $\sum w_{\mathbf{i}} = 1$; in that case,

$$var(P_{\mathbf{i}}) = q^2 \sum (w_{\mathbf{k}})^2 \approx \frac{q^2}{4\pi\sigma^2}, \tag{11}$$

where q is the per pixel RMS noise. We will assume at the beginning that this is constant–ie that we are dealing with images sufficiently faint that the read and sky noise dominate the shot noise from the object. We will deal approximately with the photon-noise dominated case later. We consider the errors in one dimension at a time, and assume that the center has been arrived at by a single quartic interpolation. The full two-dimensional analysis is complex, and the Monte Carlo simulations indicate that this simple treatment is adequate; in any case, most of the region of parameter space of interest for astrometry involves images which are photon-noise limited, and this case is treated only very approximately here.

The P_k are, of course, *not* independent;

$$cov(P_i, P_k) = \frac{q^2}{4\pi\sigma^2} \exp\left(\frac{-(i-k)^2}{4\sigma^2}\right) \approx \frac{q^2}{4\pi\sigma^2}\left(1 - \frac{(i-k)^2}{4\sigma^2} + \frac{(i-k)^4}{32\sigma^2}\right) \tag{12}$$

Thus the variances of the key quantities (reverting to the notation of the last section) are

$$var(s) = var((v_+ - v_-)/2) = \frac{q^2}{8\pi\sigma^2}\left(1 - \exp(-1/\sigma^2)\right) \tag{13}$$

and

$$var(d_2) \approx \frac{3q^2}{16\pi\sigma^6}. \tag{14}$$

to fourth order. It is interesting that the variance of d_2 vanishes in second order; the curvature in some sense goes away more rapidly than the slopes when random data are smoothed, which, I suppose, makes sense. It is also useful and interesting to note that s and $d2$ are independent random variables.

Since near the center of a star, $< d_2 > \approx A/(2\sigma^2)$, the RMS fractional error in the curvature is

$$\frac{\sigma(d_2)}{< d_2 >} \approx \frac{0.48q}{\sigma A}. \tag{15}$$

For the astrometric stars, with $a \geq \sim 300$ electrons, $q \approx 19$ electrons (including 13 electrons read noise, 30 electrons background, and an approximate contribution from shot noise from the *star*) the fractional RMS error from uncertainties in d_2 is about 7 percent, or about 7 milliarcseconds at worst (for $x \approx .25$ pixel). The standard deviation of s/d_2, neglecting the variance in d_2, is

$$\sigma(s/d_2) = \frac{q\sigma}{A\sqrt{2\pi}}\sqrt{1 - \exp(-1/\sigma^2)}, \tag{16}$$

or, for $\sigma \approx 1$,

$$\sigma(s/d_2) \approx 0.32(q/A) \approx 0.71(q/a). \tag{17}$$

which for the same star is about .045 pixels, 18 milliarcseconds. The variance in d_2 increases this to a bit more than 19 at worst, so it is errors in the slope which dominate the noise in the quadratic centering estimate. Since s and d_2 are independent random variables, this root-sum-square approach is OK.

The errors with the quadratic terms added increase, of course. With the same approximation of neglecting the variance in d_2, it is clear from

Equation 7 that, for small x, the standard deviation is increased by a factor of $1+k/(8\sigma^2) \approx 1.16$; at larger x, the factor decreases because the multiplier $(1-4x^2)$ decreases. If the variance in x is small compared to unity, the standard deviation $s4$ of the 4th-order estimate is

$$s4 = s2\left(1+\frac{k(1-12x^2)}{8\sigma^2}\right) \approx s2\left(1+0.16(1-12x^2)\right). \tag{18}$$

where $s2$ is the standard deviation of the second-order estimate. Qualitatively, this is what is observed, but higher-order effects cause a variation of the variance of $s2$ with location in the pixel which partly cancel the variation predicted here, and, in fact, the variation of the standard deviation with location in the pixel is reasonably small, as can be seen from the accompanying tables. The agreement with these theoretical values is also excellent. This is not a tribute simply to the correctness of the arithmetic, but vindicates some of our admittedly dubious expansions and even more dubious treatment of photon noise, rationalized below.

2.4 The Effects of Photon Noise

To the extent that photon noise in an object contributes significantly to the noise introduced by read noise and shot noise in the background, the pixel standard deviation k in the original image becomes a function of position in the image, and the noise structure in the convolved image becomes very complex indeed. For the estimates used above, we have simply used for q the expression $q_{eff} = \sqrt{a_{eff} + RN^2 + bkg}$, where RN (here taken to be 13 electrons) is the read noise, bkg is the sky+dark background, here taken to be 30 electrons (these values are appropriate to the astrometric chips), and a_{eff} is the effective number of electrons in a pixel near the center of the star, here taken to be the mean of the nine pixels centered on the center of a pixel-centered star image, about $0.53a$ for $\sigma \approx 1$. This simplistic approach neglects many complications, but seems to work well. All astrometrically interesting stars are well into this shot–noise–in–the–object–limited regime.

2.5 Complications at the Faint End

The second difference d_2 appears in the denominator of both the second- order estimate for x and the fourth-order correction. For the limiting bridge star considered with 300 electrons peak signal in the original image, the standard deviation of d_2 is a factor of 14 smaller than its expectation value, and one need never worry about noise causing d_2 to approach dangerously close to zero. For objects near the survey limit, however, this is not the case. The detection limit in the photometric pipeline is at 5σ and is always background limited in the photometric chips. We have seen that the total amplitude of a star in the smoothed frame is about 0.55 of its amplitude in the original, and the noise is smaller than that in the original by a factor of $\sigma\sqrt{4\pi}$. Thus if the pixel noise in the original is q and $\sigma \approx 1$, the amplitude in the smoothed frame at the detection limit is $A \approx 5q/(\sigma\sqrt{4\pi}) \approx 1.4q$. From Equation 15, the ratio of the standard deviation of d_2 to its mean is only about a third, and the estimate will fail spectacularly a few times in a thousand tries. Let us investigate how it might fail.

First consider the case in which the smoothed 3×3 array is such that the central pixel is a *real* maximum; i.e. the pixel values at all of its neighbors are strictly less than its value. Then the d_2's (let us call them d_{2x} and d_{2y} in the x and y directions) clearly cannot be zero, and, in fact, the absolute value of the offset x (and y in the orthogonal direction) must be less than 1/2. If we relax the absolute maximum condition, but demand that at least one pixel in the x direction and one in the y direction be strictly smaller than the central value, the same considerations apply except that now the x and y values can assume the extremes, ± 0.5. The absence of this condition might occur due to noise, in which case one might have a plateau of values all the same at the maximum, or at higher signal-to-noise, a very large object might cause this condition. The vanishing of the central d_{2x} or d_{2y} already invalidates the simple quadratic estimate, but for the 2-d quartic one to work all *six* of the second differences associated with the 3×3 array must be positive. In any of the cases in which this fails to be true the center is not very well defined anyhow, so exactly what one does is probably not too important, so long as the centering algorithm fails gracefully. We describe one possible approach in the next section.

2.6 Computation

The procedure to be followed is outlined here, and the first attachment contains a set of C functions which produce the centroids by these techniques as they might be implemented in the astrometric DA or other case in which one does *not* have the whole smoothed image to search for the maximum. They are not well optimized for speed but will produce a centroid in of order 1 millisecond per image on a 15-20 MIPS machine with a modest cache, such as the DA processor. The procedure is as follows:

1. Find a maximum in the smoothed image. This is procedurally different for the case in which one has the entire smoothed image at one's disposal, as in the photometric pipeline, and not, as in the astrometric DA. In either case, if there is a connected region three pixels or larger in extent (i.e., which span a range of three or more pixels in either x or y) in which the value of the smoothed image is at the maximum, the position is taken to be at the centroid of this connected patch. This is probably not very good, but the position in this case is ill-defined at best and this will do and is definite. For the astrometric DA, these cases all are indicative that the object is *not* a star and is to be rejected. Notice that this procedure insures that both d_2x and d_{2y} for the central crosstripe are positive.

2. Find the simple quadratic position offsets from the center using

$$x_0 = \frac{s_x}{d_{2x}}$$

$$y_0 = \frac{s_y}{d_{2y}},$$

where s_x and s_y, are, of course, the slopes in the relevant directions. They are determined for this first estimate along the centered cross; i.e..

$$x_0 = \frac{P_{i,j+1} - P_{i,j-1}}{2(2P_{i,j} - P_{i,j-1} - P_{i,j+1})}$$

$$y_0 = \frac{P_{i+1,j} - P_{i-1,j}}{2(2P_{i,j} - P_{i-1,j} - P_{i+1,j})}. \tag{19}$$

We do this to have a starting point for the iteration for the real maximum later and to have a "pretty good" position to return in case the full fourth-order treatment fails, which it can still do at this point.

3. Evaluate the six one-dimensional fourth-order maxima along the rows and columns of the 3×3 array. If any of the six second differences vanishes or is negative, return the quadratic position estimate. The three maxima along the rows (for example) typically lie along an approximately vertical,

almost straight line. The maxima define the locus of horizontal maxima in the image, which we approximate by a parabola:

$$mx(y) \approx mx_0 + y(mx_+ + mx_-)/2 - y^2(2mx_0 - mx_+ - mx_-),$$

with a similar expression for the locus of the vertical maxima my. The sought-for maximum is at the intersection of those loci, $(mx(y), y) = (x, my(x))$. Approximate mx using the quadratic estimate for my, and my using the quadratic estimate for mx. Iterate once.

4. Return.

2.7 Computational Results

A series of Monte Carlo computations have been run using the above scheme. The procedure was as follows: A single image from the astrometric array, centered on a pixel, as produced with a full raytrace of the system and smoothed with the standard seeing model, was replicated 2500 times in a 50x50 regular array separated by 20 pixels. This image is sinc-shifted as a whole to produce offsets from centration on a pixel. Photon and read noise was added and the center of each image found by the above procedure. Statistics were done on the results, yielding mean offsets and standard deviations for the estimated centroids. The results are in good agreement with the analytic estimates; they are presented in the accompanying Table 2.

As perusal of this table shows, the results are quite good, with systematic errors (as exemplified by the high signal-to-noise cases) of the order of 3 mas or less everywhere, and the offsets are in good agreement with the noiseless, pixel-centered offsets between the real centroids and the gaussian-least-squared fits (to within a fraction of a mas). The standard deviations are quite uniform over the array of pixel offsets and in good agreement with the estimates above. The standard deviations for the good images (all the detectors at all field positions except chips 18 and 21) at central fluxes of 300 electrons are all about 20 mas, using a smoothing gaussian of sigma 1.2 pixels, compared to 19 as estimated above for $\sigma = 1$. This degrades to about 22 mas at the edge of chip 21 and 24 mas at the edge of field 18 for a star of the same total flux. The latter field position has the worst images in the array (effective $\sigma \approx 1.4$; we used one smoothing σ for the whole array which was a compromise value, though there is no reason to do this in practice).

We will return to the issue of how best to choose the smoothing sigma in a subsequent paper, which will also deal with how to focus the telescope using related techniques.

3 Saturated Images

Photo finds the centre of saturated images by considering the marginal distribution of image intensity in x and y separately. Specifically, it takes the bounding box provided by the object finder, grows it by 3 pixels in each direction, rejects any columns that contain saturated pixels, and projects the remaining pixels onto the x- and y-axes.

In order to find the centre, it fits a model of the PSF to the two marginal distributions. In 2-D, the model adopted is very similar to that in the Grey Book:

$$P(r) = \alpha e^{-r^2/2\sigma_1^2} + (1-\alpha)e^{-r^2/2\sigma_2^2} + f\left(1 + \frac{r^2}{3\sigma_2^2}\right)^{-3/2}. \tag{1}$$

Integrating this expression over all r, the power in the three components is $2\pi\alpha\sigma_1^2$, $2\pi(1-\alpha)\sigma_2^2$, and $6\pi f\sigma_2^2$ respectively.

We may then project P onto a row, giving

$$P_x(x) = \sqrt{2\pi}\alpha\sigma_1 e^{-x^2/2\sigma_1^2} + \sqrt{2\pi}(1-\alpha)\sigma_2 e^{-x^2/2\sigma_2^2} + 2\sqrt{3}f\sigma_2\left(1 + \frac{x^2}{3\sigma_2^2}\right)^{-1} \tag{2}$$

and

$$\begin{aligned} P_y(y) = {} & \sqrt{2\pi}\alpha\sigma_1 e^{-y^2/2\sigma_1^2}[\mathrm{erf}(\xi/\sqrt{2}\sigma_1)/2] + \\ & \sqrt{2\pi}(1-\alpha)\sigma_2 e^{-y^2/2\sigma_2^2}[\mathrm{erf}(\xi/\sqrt{2}\sigma_2)/2] + \\ & \sqrt{3}f\sigma_2\left(1 + \frac{y^2}{3\sigma_2^2}\right)^{-1}\left(1 - \sqrt{\frac{\xi^2/3\sigma_2^2}{1 + (y^2+\xi^2)/3\sigma_2^2}}\right) \end{aligned} \tag{3}$$

where ξ is the distance from the image's centre of the first good column; in general we have to add the left and right components of P_y to arrive at the total marginal distribution. As expected, if $\xi = 0$, adding the two parts of P_y recovers the functional form of P_x.

For centroiding, we take $\sigma_2 = 3\sigma_1$, and assume that the two Gaussian components have equal power; this leads to

$$\alpha = 0.9$$

$$f = \frac{\lambda}{15}$$

where λ is the fraction of the total power in the power law wings. We then set $\sigma = 1$; corresponding to seeing of about 1 arcsec.

It will be noted that the last term in equation 3 is not a pure power law; rather than face this fact properly, we note that for that part of the marginal distribution that the power-law wings are important, $y \gg \xi$, so the term in the square root is small. As the exact form of the profile is not crucial to the centroiding, we simply set y to some characteristic value, and render it a constant.

Once we have the marginal distributions of the data and of the model, the x and y positions are derived independently using Brent's algorithm to minimise the mean-square deviation; for each estimate of the centre, the amplitude of the PSF may be calculated as the ratio of the sum of the data to that of the model.