Flux-Uncertainty from Aperture Photometry

F. Masci, version 1.0, 10/14/2008

1. Summary

We derive a general formula for the noise variance in the flux of a source estimated from aperture photometry. The 1-σ uncertainty is given by the square root of this expression. Unless mentioned otherwise, it is assumed that the flux estimate includes background subtraction and that there are *no severe pixel-to-pixel correlations*. These usually occur in resampled and interpolated images (e.g., mosaics). If present, the uncertainty computed by the formula below will be an underestimate. Correlated noise is rather complicated to express analytically. A formalism to account for correlated noise using Monte Carlo simulations is described in:
http://web.ipac.caltech.edu/staff/fmasci/home/wise/ApPhotUncert_corr.pdf

The only difficult thing to determine before using this formula is the gain factor *g*, i.e., how many electrons correspond to a pixel flux unit if the image at hand were *directly* observed with the detector. Electrons are the discrete entities that are counted, and these contribute to Poisson fluctuations. A co-add or mosaic of images is *not* directly observed. These are likely to be resampled to a different pixel size, and are typically generated from an average (or median) of a number of overlapping detector images. Co-addition has the effect of reducing the Poisson noise-uncertainty by $1/\sqrt{N_i}$, where N_i is the depth-of-coverage at image pixel i. N_i is usually constant over a source, and can be approximated as the average or median depth over all pixels in the source aperture. All other noise contributions in the formula below (other than the Poisson term) can be derived *a posteriori* from the image pixels at hand.

Here is the final result. We also define all quantities involved. The derivation is given below.

$$\sigma_{src}^2 = \frac{1}{g}\sum_{i}^{N_A}\frac{\left(S_i - \overline{B}\right)}{N_i} + \left(N_A + k\frac{N_A^2}{N_B}\right)\sigma_{B/pix}^2$$

where

g = gain in "electrons / pixel data units" pretending that image was "observed" by a detector with the same pixel size. Note, the image at hand could be a resampled co-add or mosaic

N_A = number of pixels in source aperture

N_B = number of pixels in background annulus

N_i = depth-of-coverage at pixel i; $N_i > 1$ if image is a co-add

S_i = signal in pixel i in image data units

$\overline{B}$ = estimated background per pixel in annulus (either mean or median):

if $\overline{B}$ = mean background/pixel, $k = 1$

if $\overline{B}$ = median background/pixel, $k = \pi/2$

if assume $\overline{B} = 0$ or if no background is subtracted, set $k = 0$

$\sigma^2_{B/pix}$ = variance in sky background annulus in [image units]^2/pixel.

Can compute from square of RMS deviation from mean or median. Can also approximate using a robust estimator of scale:

$\approx [0.5(q_{0.84} - q_{0.16})]^2 \approx [(q_{0.5} - q_{0.16})]^2$ where the q are quantiles;

This last approximation is even more robust since it only uses the lower tail where cosmic rays and spurious sources are less likely to occur

N.B : if first summation term above is < 0 (due to $\overline{B}$ exceeding source photon fluctuations in aperture, set this term to zero. It means the measurement is consistent with zero Poisson noise from the source.

2. Derivation

The above formula is derived as follows. First, the equation for estimating the flux of a source from aperture photometry can be written:

$$F_{src} = F_{tot} - N_A \overline{B}, \tag{1}$$

where F_{tot} is the total integrated flux in the aperture, and other quantities are defined above. The noise-variance in this estimate is given by standard error propagation. We ignore correlations between pixels in the source aperture and background annulus since these are assumed to be well separated. The variance in Eq. (1) is given by:

$$\sigma^2_{src} = \sigma^2_{tot} + N_A^2 \sigma^2_{\overline{B}}, \tag{2}$$

where the line above the sky background B indicates that its estimate refers to either a mean or median value.

The first term on the right hand side of Eq. (2) is the total variance in the source aperture. This is the quadrature sum of all the individual pixel variances therein:

$$\sigma_{tot}^2 = \sum_i^{N_A} \sigma_i^2 \; . \tag{3}$$

For simplicity, we ignore possible pixel-to-pixel correlations. These will be handled in a future revision.

An individual pixel variance σ^2_i (for any pixel in the source aperture) can be written in terms of a Poisson variance term to account for the contribution of photon noise to pixel i from the actual source, and a term to account for all other "extraneous" noise components, e.g., photon-noise from the sky, read-noise and other instrumental noise. Recall that the source photons are superimposed on a background, and we must account for the contribution from each noise component separately. In general, if the image at hand were *observed* with the detector (with the same pixel size), the noise variance in units of *electrons*^{2} can be written:

$$\sigma_{ie}^2 = g\left(S_i - \overline{B}\right) + \sigma_{Be}^2, \tag{4}$$

where the $(S_i - \overline{B})$ is the signal in a pixel from source alone, in physical units of the image data, g is the gain in "*electrons*/*pixel data units*", and σ^2_{Be} is the extraneous noise component in *electrons*^{2} (see above). The pixel-variance in (4) can be written in units of [*image units*]^2 by dividing through by g^2:

$$\sigma_i^2 = \frac{\left(S_i - \overline{B}\right)}{g} + \sigma_{B/pix}^2, \tag{5}$$

where $\sigma^2_{B/pix}$ is now the extraneous "background" noise-variance in physical image units. If the image at hand is a co-add made from a stack of overlapping images, we must divide the Poisson source term by the number of overlaps at pixel i: N_i. Note, we do not divide the background variance term. This is because this term will be computed directly off the co-add pixels (e.g., via a spatial pixel RMS or data scale), and this implicitly includes the $1/N_i$ variance-reduction. Poisson noise from the source cannot be estimated in a similar manner because the source does not occupy a large enough region to compute a spatial RMS. Furthermore, the source is likely to have a profile that varies rapidly over pixels. It is difficult to isolate the Poisson fluctuations "about this profile". It can be estimated from the variance in the residuals of a profile fit, but this is outside the scope of this document. Therefore, to predict the amount of Poisson noise from the source, we must go back to the detector image-frame where the "electron counting" is done. Equation (5) can then be re-written:

$$\sigma_i^2 = \frac{\left(S_i - \overline{B}\right)}{gN_i} + \sigma_{B/pix}^2, \tag{6}$$

Combining Equations (2), (3) and (6), the noise-variance in a background-subtracted source flux is given by:

$$\sigma_{src}^2 = \frac{1}{g}\sum_i^{N_A}\frac{\left(S_i - \overline{B}\right)}{N_i} + \sum_i^{N_A}\sigma_{B/pix}^2 + N_A^2\sigma_{\overline{B}}^2. \tag{7}$$

Without prior knowledge of the variation of the "background" pixel variance $\sigma^2_{B/pix}$, we simply assume it is a constant, i.e., independent of pixel location. This is computed from pixels in the sky annulus. Therefore, we can replace the second summation term with the number of pixels in the source aperture, N_A. Equation (7) now becomes:

$$\sigma_{src}^2 = \frac{1}{g}\sum_i^{N_A}\frac{\left(S_i - \overline{B}\right)}{N_i} + N_A\sigma_{B/pix}^2 + N_A^2\sigma_{\overline{B}}^2. \tag{8}$$

Now to the last term – the variance *in the* mean or median sky-background per pixel. If $\overline{B}$ is an arithmetic mean, its variance (or the square of its 1-sigma uncertainty) is given by

$$\sigma_{\overline{B}=\mu}^2 = \frac{\sigma_{B/pix}^2}{N_B}, \tag{9}$$

where $\sigma^2_{B/pix}$ is our background pixel variance from above, computed using the N_B pixels in our sky annulus.

If however the median was used for $\overline{B}$ and the pixel values are normally distributed, its 1-sigma uncertainty as derived from Eq. (9) will be slightly underestimated by a factor of $[\pi/2]^{1/2}$, or its variance underestimated by $\pi/2$. In other words, the median is noisier (*less efficient* in statistical parlance) than the mean for a randomly drawn sample. Nonetheless, given the robustness of the median against outliers, this is a small price to pay. A derivation of this "$\pi/2$ inflation" exists in each of the following references and was used to derive Equation (8) in the following paper:

http://web.ipac.caltech.edu/staff/fmasci/home/statistics_refs/MADstats.pdf

[2] B. L. van der Waerden, Mathematical Statistics, Springer, New York, 1969, section 17.
[3] S. S. Wilks, Mathematical Statistics, Wiley, New York, 1962, section 9.6.

Therefore, under the assumption of normally distributed data (which is usually satisfied in the limit of large N_B with a 'well behaved' astronomical detector), the variance *in the* median is given by:

$$\sigma_{\overline{B}=med}^2 = \frac{\pi}{2}\frac{\sigma_{B/pix}^2}{N_B}, \tag{10}$$

We can combine Equations (9) and (10) with Equation (8):

$$\sigma_{src}^2 = \frac{1}{g}\sum_i^{N_A}\frac{\left(S_i - \overline{B}\right)}{N_i} + \left(N_A + k\frac{N_A^2}{N_B}\right)\sigma_{B/pix}^2. \tag{11}$$

Where $k = 1$ corresponds to $\overline{B}$ estimated using an arithmetic mean, and $k = \pi/2$ is for $\overline{B}$ estimated using a median. Incidentally, if the source flux estimate involved no sky-background subtraction, or the background is known to be negligible *a priori*, one can set $k = 0$.

One can use robust estimators of spread in the sky-pixel values as a proxy for $\sigma^2_{B/pix}$. This is to avoid estimates (like the RMS) from being biased by cosmic rays and spurious sources in the sky annulus. Examples include using the quantile difference: $\sigma_{B/pix} \approx 0.5(q_{0.84} - q_{0.16})$, or the Median Absolute Deviation (MAD): $\sigma_{B/pix} \approx 1.4826 med|p_i - med\{p_i\}|$ where p_i is the value of pixel i. Both these measures assume normally distributed data, and they converge exactly to the RMS value (standard deviation) as $N_B \rightarrow \infty$.

An issue is whether these robust proxies are as *efficient* as the RMS under the assumption of a normally distributed population when 'small' sample statistics are involved. By efficient, we mean no more variant than the RMS itself. As an aside, the RMS has the *least* variance for random sampling from a normally distributed population. A moment's thought tells us that just like the median (the 0.5 quantile), additional uncertainty will be introduced if *robust* estimators of spread are used. This implies different corrections that can be absorbed into k. These corrections can be derived using a Monte-Carlo simulation. In the end, the benefits that these robust measures provide in the presence of outliers greatly outweighs any additional uncertainty they may introduce.

One may verify Equation (11) via a simulation. For example, one can add different noise components with known variance to an image containing a single source (a point convolved with a test PSF). One then uses Equation (11) to compute the uncertainty in the source flux estimated from aperture photometry. This is then compared with the input truth.