Impact of Noisy Pixels on Photometric SNR

F. Masci, version 1.0 (3/2/2009)

Given a distribution of the pixel noise-variance derived from dark data, we ask the following: (i) what is the reduction in the average Signal-to-Noise Ratio (SNR) of a point source on average if a single pixel with some abnormally high variance is not masked prior to photometry, and (ii) what fraction of sources will have their SNR reduced by this amount. The goal is to determine a threshold for the pixel variance above which pixels should be masked. One will do better using a simulation, but the below should provide some ballpark estimates.

<u>Assumptions</u>:

- Only one noisy pixel is assumed to fall within the effective footprint of a source, i.e., as defined by the number of noise pixels for the Point Response Function (PRF). This is reasonable if the fraction of noisy pixels above some threshold is known to be low.
- The calculations below assume photometry (e.g., profile fitting) on a single frame. If the photometry is performed by combining sources from N frames, then the effective point-source flux uncertainty will be *lower* (and SNR higher) than that predicted from $\sqrt{N}$ statistics alone. This is because not all sources across the frames will fall on a bad (noisy) pixel. The contribution from the "good" $N - p$ sources will dilute the effect of the p source(s) containing the bad noisy pixel(s). Therefore, the below can be seen as a 'worse-case' scenario. One will always do better when redundant measurements are combined.
- True profile-fit photometry usually involves χ^2 minimization where the procedure implicitly includes inverse-variance weighting using priors for the pixel variances. This effectively reduces the impact of a noisy pixel on the both the estimated flux and uncertainty. The flux uncertainty as predicted below (in the denominator of the SNR expression) is expected to be slightly larger than that derived from a weighted fit. This also leads to a worse-case scenario for the below.

The point-source SNR as a function of the factor x by which a single pixel is deviant with respect to some nominal 'dark' pixel sigma σ_{good} (e.g., the modal or median sigma) is defined as:

$$SNR(x) = \frac{N_p S_p}{\sqrt{N_p S_p + g^2 x^2 \sigma_{good}^2 + (N_p - 1) g^2 \sigma_{good}^2 + N_p \sigma_{sky}^2 + N_p \sigma_{other}^2}},$$

where:

N_p = effective number of 'noise pixels' for the PRF in question;
S_p = average point source signal *per pixel* to assume over region covered by N_p, in e-. Instead of assuming explicit values, we pick a range of nominal SNR(x=1) values (e.g., 5,10,15…100) and invert the above formula to compute the required S_p. This is then used for computing the general SNR(x).
g = gain in e-/(DEB DN);
σ^2_{good} = nominal dark pixel variance in DN^2 per pixel, assume = mode of pixel noise-variance distribution;
x = factor by which a single bad (noisy) pixel is deviant relative to σ_{good}, i.e., $x = \sigma_{bad} / \sigma_{good}$ where σ_{bad} is some threshold > σ_{good} in the pixel-sigma distribution;
σ^2_{sky} = photon noise-variance in (e-)^2 per pixel from sky background alone, = $S_{sky/pix}$ in e-. We assume average values for $S_{sky/pix}$ as expected for WISE over all ecliptic latitudes;

σ^2_{other} = all other contributions to the pixel noise-variance (ignored at present).

Values assumed for the above parameters are summarized below.

	source	W1	W2	W3	W4
N_p [native pix]	1	13.5	16.6	37.1	27.0
g [e-/DN]	3	3.827	3.827	4.725	4.725
σ^2_{good} [DN^2]	4	9	9	289	90
σ^2_{sky} [e-^{2}]	5	14.7	95.6	6028.6	6224.6
σ^2_{other} [e-^{2}]	2	0	0	0	0

Sources:

1. Mark Larsen (SDL), June 2008
2. Assumption
3. Mark Larsen (SDL), Dec 2008
4. Dark data from flight-model MIC2 testing, Nov 2008. Characterized by M. Skrutskie
5. WISE Calibration Plan

Below we show plots of the quantity:

$$R(x) = \frac{SNR(x)}{SNR(x=1)}$$

as a function of x (= $\sigma_{bad} / \sigma_{good}$) for each band. This quantifies the reduction in SNR relative to the nominal case SNR(x=1) when a single 'bad' pixel with sigma $x\sigma_{good}$ contributes to the flux uncertainty. Curves are shown for a range of nominal SNR(x=1) values: from bottom to top, SNR(x=1) = 0.5…100 in steps of 5. As expected, the impact of a single bad (noisy) pixel at high values of SNR (where source-flux dominates) is considerably reduced. The impact is also reduced when the sky-background is high, as in bands 3 and 4.

The fraction of sources which will have their SNR's reduced by some $R(x)$ or more on average is effectively given by the fraction of pixels in the noise-sigma distribution with $\sigma > x\sigma_{good}$.

How do we interpret the plots below? Take for example **band 1**. Let's adopt a noise-sigma threshold of x = 10. This corresponds to pixels with $\sigma > x\sigma_{good}$ = 10*3 = 30, where σ_{good} was taken as the mode. This means that sources with nominal SNR ~ 5 (i.e., the SNR they would have had if all was perfect) would be reduced to SNR ~ 0.4*5 = 2. Note that the nominal SNR ~ 5 curve is the second from the bottom. Similarly, sources with nominal SNR ~ 20 will be brought down to SNR ~ 20 * 0.45 = 9 (5th curve from the bottom). From the band 1 pixel sigma distribution, ~0.2% of the pixels have x > 10 (from M. Skrutskie). This means that ~0.2% of sources can potentially have their SNRs reduced by at least the above-mentioned amounts, assuming of course they're affected by only one bad (noisy) pixel. This fraction can definitely be tolerated since the WISE completeness requirement asks that no more than 5% of sources with nominal SNR > 20 be unaccounted for.

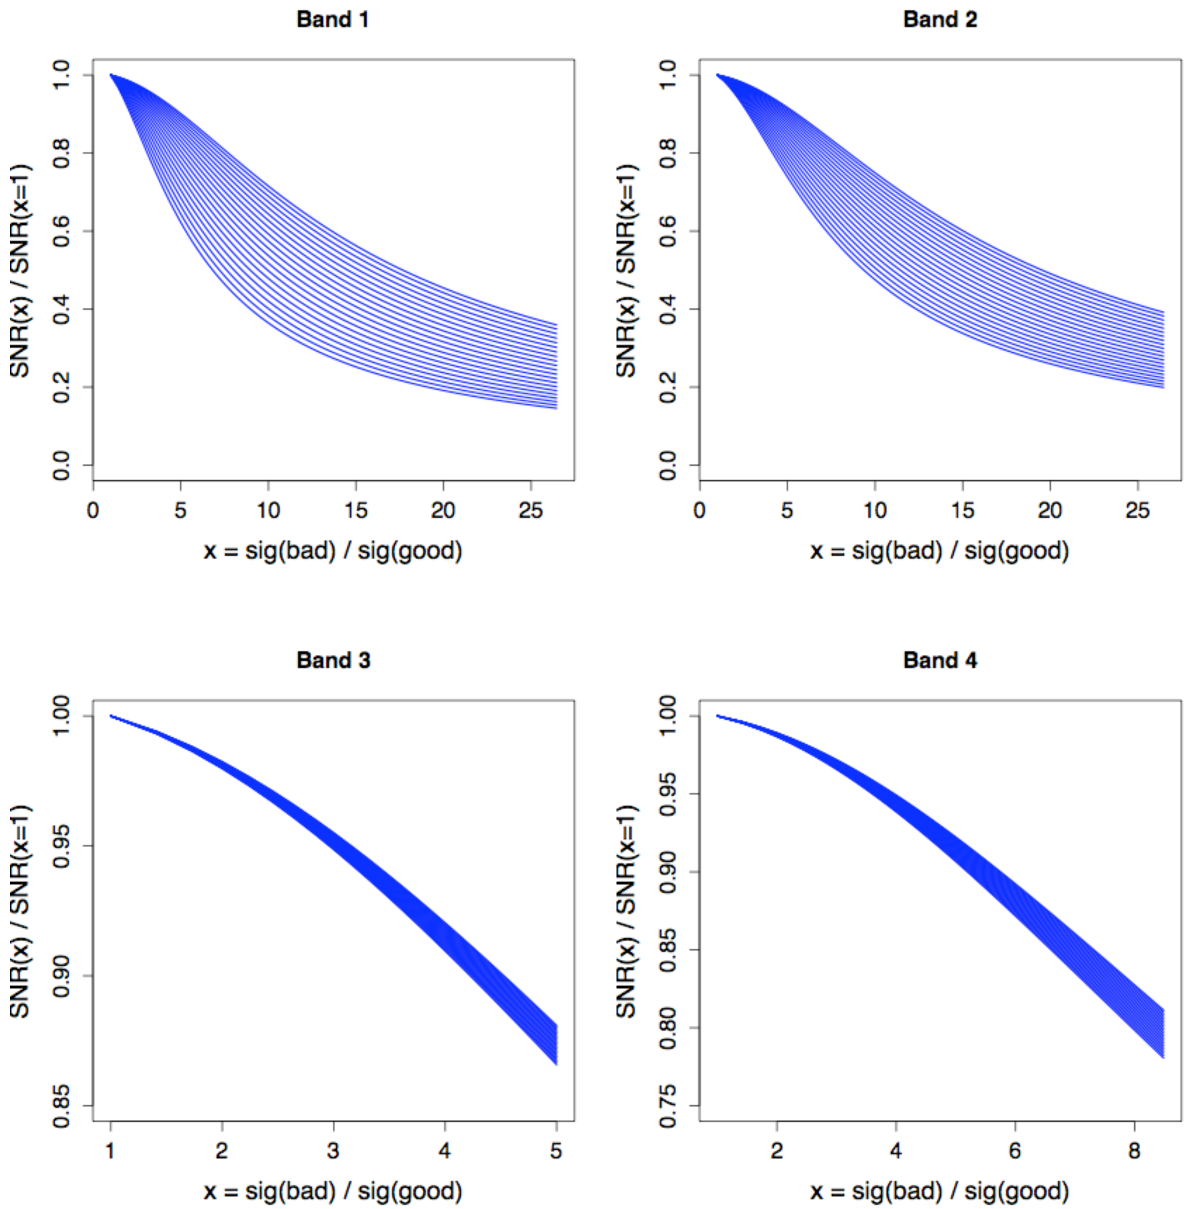
Band 1
Band 2
Band 3
Band 4
SNR(x) / SNR(x=1)
x = sig(bad) / sig(good)