Image Blurring and Convolution, and Pixel Correlation and Covariance

June 15, 2009

John W. Fowler

An image composed of N_c columns and N_r rows contains $N_c \times N_r$ pixels P_{ij}, $i = 1$ to N_c and $j = 1$ to N_r. The pixel matrix can be represented as a vector composed of elements p_n, $n = 1$ to $N_t = N_c \times N_r$, where

$$\begin{aligned} p_n &\equiv P_{ij} \\ n &\equiv i + (j-1) N_c \end{aligned} \tag{1}$$

We consider an image of flux values F_n such that all pixels are uncorrelated, i.e., any blurring of the fluxes is completely undersampled by the pixel sizes, and we consider each flux to be a random variable independent of all other fluxes in the image. We define the corresponding image containing pixels f_n to be the result of subtracting each mean $\langle F_n \rangle$ from F_n, so that each f_n is a zero-mean random variable with some non-zero variance that represents our ignorance of $\langle F_n \rangle$, which represents the true value of the pixel.

We now "observe" the image with an instrument whose response function is oversampled by the pixel size. This results in a data image of pixels D_m, where each pixel also has a mean and additive random variable d_m. The observational process maps the F image to the D image in a linear manner that conserves flux. The response function that produces D_m as a linear combination of pixels F_n is R_{mn}, i.e.,

$$D_m = \sum_n R_{mn} F_n \tag{2}$$

Because the mapping is linear, subtracting off the means results in

$$\begin{aligned} d_m &= \sum_n R_{mn} f_n \\ \vec{d} &= R \vec{f} \\ \vec{f} &= R^{-1} \vec{d} \end{aligned} \tag{3}$$

Note that, other than being normalized to conserve flux, R_{mn} is a completely general matrix that maps fluxes in a manner that causes them to be blurred, because it describes a response function that is oversampled by the image pixels. In the special case wherein the shape of the response function is independent of location within the image, Equation 3 reduces to a discrete convolution, and the response function is called "isoplanatic". We do not make this restriction, and so we use the term "blurring" to indicate the more general (non-isoplanatic) mapping.

We consider two distinct data pixels, d_i and d_j, and compute their covariance.

$$
\begin{aligned}
d_i &= \sum_m R_{im} f_m \\
d_j &= \sum_n R_{jn} f_n \\
\operatorname{cov}(d_i, d_j) &= \left\langle \left(\sum_m R_{im} f_m \right) \left(\sum_n R_{jn} f_n \right) \right\rangle \\
&= \left\langle \sum_{m,n} R_{im} R_{jn} f_m f_n \right\rangle \\
&= \sum_{m,n} R_{im} R_{jn} \left\langle f_m f_n \right\rangle
\end{aligned} \tag{4}
$$

Since the *f* pixels are uncorrelated, the only terms that survive the summation are those for $m = n$, and therefore

$$
\begin{aligned}
\operatorname{cov}(d_i, d_j) &= \sum_n R_{in} R_{jn} \left\langle f_n^2 \right\rangle \\
&= \sum_n R_{in} R_{jn} \sigma_{fn}^2
\end{aligned} \tag{5}
$$

The corresponding correlation coefficient depends on the variances of d_i and d_j. Since

$$
\sigma_{dn}^2 = \left\langle d_n^2 \right\rangle = \left\langle \left(\sum_m R_{nm} f_m \right)^2 \right\rangle = \sum_m R_{nm}^2 \left\langle f_m^2 \right\rangle = \sum_m R_{nm}^2 \sigma_{fm}^2 \tag{6}
$$

we have

$$
\rho_{ij} = \frac{\sum_n R_{in} R_{jn} \sigma_{fn}^2}{\sqrt{\left(\sum_n R_{in}^2 \sigma_{fn}^2 \right) \left(\sum_n R_{jn}^2 \sigma_{fn}^2 \right)}} \tag{7}
$$

In principle, the correlation can be removed by diagonalizing the covariance matrix. These diagonalized elements are the eigenvalues of the covariance matrix, and the corresponding eigenvectors define a rotation in N_t-dimensional space which could be applied to the *d* image to obtain its corresponding uncorrelated values. But since those constitute the *f* image, that mapping is already given by the last line in Equation 3, and so the inverse of the response matrix is the matrix that is needed. Note that this is not generally the same as the rotation matrix that diagonalizes the covariance matrix, however, because in general the *R* matrix is not orthonormal and not even square if the pixel sizes in the two images are different or if multiple data images are derived from the *F* image (e.g., overlapping dithered coverage). In the general case, the inverse of R must be computed as a Moore-Penrose pseudo-inverse.

The fact that the matrix which maps F to D is not unique may be seen in Equation 2. If n runs from 1 to N_t and m runs from 1 to M_t, and assuming that D_n and F_m are known, then Equation 2 is one equation in N_t unknowns, and we have M_t such equations. If $N_t > 1$, there are an infinite number of solutions for the R_{mn} values. If R is square and orthonormal, then it must be the same as the rotation matrix whose inverse diagonalizes the covariance matrix, but in general it is neither. In practice, R can be constructed from the response function for the observing instrument. Given D, it is probably more efficient computationally to compute F from D via the Moore-Penrose pseudo-inverse of R than to compute the covariance matrix, diagonalize it, and extract the rotation matrix from the eigenvectors. The latter procedure would yield the vector of uncorrelated elements corresponding to D, and an un-blurred image would still have to be constructed from them, presumably by co-adding with a delta-function interpolation kernel.

It must be recognized that current computer technology is not capable of doing any of these computations for astronomical images of realistic size. For example, with a single D image and an F image each of 1024 rows and columns, the R matrix contains $1024^4 = 1{,}099{,}511{,}627{,}776$ elements. Manipulating such a matrix would probably require double precision, and so the matrix would occupy almost 8.8 terabytes of memory.