Astronomical Data Analysis Software and Systems IX
ASP Conference Series, Vol. 216, 2000
N. Manset, C. Veillet, and D. Crabtree, eds.

Dithering, Sampling and Image Reconstruction

R. N. Hook

ST-ECF, Karl-Schwarzschild Str.2, D-85748 Garching bei München, Germany

A. S. Fruchter

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218, USA

Abstract. Many astronomical cameras (and spectrographs) are now routinely operated in a dithered mode – with small shifts between multiple exposures. Such dithering offers many advantages, particularly for the case in which the original detector pixels are too large to adequately sample the point-spread function of the image falling on them, but can pose a difficult image combination problem. We will describe the factors which affect the information content of dithered images and their characteristics.

There are now several methods available for reconstructing a "super image" of the sky from a set of such dithered images. The different algorithms, including the "Drizzling" method developed by the authors, the method proposed recently by Tod Lauer, and more conventional approaches based on image interpolation will be compared and contrasted and examples of their use on real data given.

Finally some comments about future developments, both of cameras and algorithms will be made.

1. Introduction

Many astronomical observations involve multiple exposures of the same field. If there are small shifts introduced in the pointing of the telescope between exposures the images are said to be "dithered". Such shifts may be deliberately introduced in a selected pattern or they may be a result of the inability of the telescope to exactly duplicate a specified pointing. The reduction of the data will generally include a step in which such dithered frames are "coadded" to produce a single output frame of improved signal-to-noise ratio.

Many detectors in astronomy have pixels which are too large to adequately sample the point-spread function of the image falling on them. Camera design is always a compromise and larger pixels give a larger field for a given number of pixels and have other advantages. Under-sampled cameras are common and a famous example is the Wide Field Planetary Camera 2 (WFPC2) on the Hubble Space Telescope.

The dithering and undersampling of astronomical cameras together pose important problems of observation planning, execution and subsequent data reduction. This paper introduces the subject and describes several methods for the combination of such data. Particular emphasis is given to the processing of data from the optical and near-IR cameras on the Hubble Space Telescope (HST). We first describe how the optics, detector and dithering affect the resultant image. The bulk of the paper is a discussion of several different image reconstruction methods and their relative merits. Some concluding remarks about the effects of the pixel-response function are also included. A practical rather than theoretical approach to these problems is adopted.

1.1. Why, and Why not, Dither?

Deliberate dithering strategies are adopted for a variety of reasons. Firstly dithering results in a given object being imaged onto different parts of the detector and hence helps with artifact detection and suppression as well as helping to compensate for flat-field effects. The latter are particularly prominent in ground-based infrared observations and "jitter" techniques are very widely used in this field (e.g. Devillard 1999). The second reason is because sub-pixel dithering can allow the reconstruction of some of the information which has been lost because of spatial undersampling. Most of the following discussion concentrates on the second of these.

Dithering is not always the optimum strategy. Executing the dithers with adequate accuracy (at a sub-pixel level) may be difficult or impose overheads. Splitting exposures may lead to additional noise (e.g., readout noise in cases where the images are detector rather than sky noise limited) and dithering inevitably leads to a smaller final field of the deepest imaging. Finally the precise measurement of shifts between images and the reconstruction of the final co-added image can be time consuming.

1.2. Aspects of Image Formation

The formation of a digital image by a typical astronomical telescope and detector (e.g., a CCD) combination involves a series of degradations and loss of information which are illustrated in Figure 1. The true intensity distribution on the sky is first convolved with the point-spread function (PSF) of the telescope and camera optics. This effect may vary with time, colour and position on the detector. When this new image falls on the detector the intensity distribution is convolved with the pixel-response function (PRF), which may also vary from pixel to pixel and with colour, and is then sampled at the centre of each pixel. If the assumption is made that the pixel is a rectangle of uniform sensitivity this is equivalent to the integration of the intensity over the pixel. The final stage is the conversion of this pixel intensity into a digital data value allowing for photon and other sources of noise. This process may be represented as

$$I(i,j) = Sky(i,j) \otimes PSF(i,j) \otimes PRF(i,j) \tag{1}$$

where $\otimes$ is the convolution operator and i, j are indices over a fine grid which well-samples the images. $I(i,j)$ is then sampled at the pixel centres of the final output detector. This sampling may be regarded as a further multiplication with a two-dimensional grid of δ-functions (a Shah function).

For an undersampled camera system the separation of the sampling points is normally the same as the width of the pixel and hence the extent of the PRF but this need not always be the case. It is important to note that in typical examples such as the WFC channels of WFPC2 the convolution with the PRF (approximately a square "box" of width 0.1") causes a loss of high spatial frequency information which is greater than that due to the convolution with the PSF (an Airy function with FWHM of about 0.05" at 500nm).

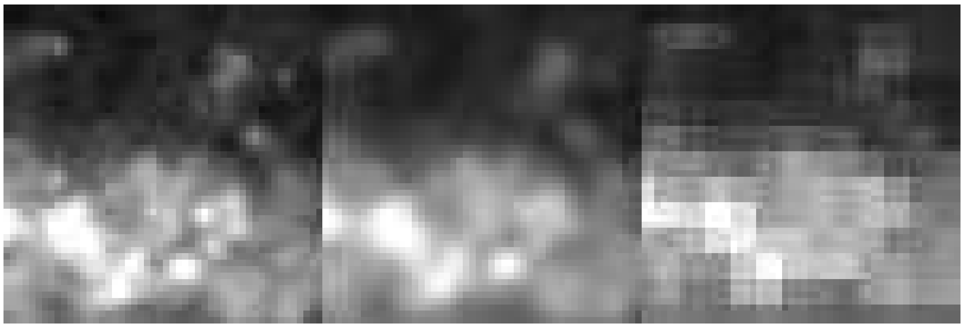

Figure 1. The image formation process. The left-hand image is the true intensity distribution on the sky convolved with the telescope point-spread function (PSF) — this is the image falling on the detector. The central image shows the result of convolving with the pixel-response function and the right-hand one the final outcome when this is sampled at the centre of each pixel.

When multiple dithered images are taken with an adequate number of sub-pixel offsets the sampling of the intensity distribution $I(i,j)$ is improved and, in principle, should allow a full reconstruction. This is the aim of most of the methods discussed below. No amount of dithering can compensate for the loss of information introduced by the convolution with the PRF but we can hope to combine a set of well-dithered images resembling that on the right in Figure 1 to produce a result similar to that at the centre.

2. The Reconstruction of Dithered Images

Over the past few years the WFPC2 and other cameras on HST have produced a vast amount of superb quality undersampled imaging data. The desire to fully exploit this data has led to considerable interest in methods for the reconstruction of optimal combinations from dithered undersampled images. This section reviews some of the most widely used methods and provides pointers to further information.

2.1. Interlacing, "Shift-and-Add" and Interpolation-Based Methods

If images can be obtained with precise fractional pixel offsets on a regular grid (typically four pointings with half-pixel shifts in both directions) then a combination can be created on a pixel grid finer than the original by interlacing the data values. Such a combination involves no resampling or additional image degradation and is optimal. Unfortunately such exact offsets are rarely possible and in addition the presence of geometric distortion means that they are only exact in a restricted region of the detector.

Another simple approach is to shift each image to bring them into alignment and then coadd the results. The coaddition can then reject anomalous values such as cosmic rays or hot pixels. This may be done either using simple "shift-and-add" or a more sophisticated interpolation method. For well-sampled data and a careful choice of interpolator (e.g., sinc) this technique can work very well and is often used on ground-based images. Fully developed software for these tasks exists and is widely used (e.g., **imcombine** in IRAF and its variants in other packages). Unfortunately this is not true of undersampled data which, when interpolated, will inevitably suffer from artifacts and smoothing which will also have the side-effect of smearing small image defects such as cosmic rays and hence making their detection and suppression less effective. Shift-and-add normally involves replacing each pixel from the input image by a square of the same size before the "shift" stage and hence involves an additional convolution with the PRF and corresponding degradation of resolution. Some implementations of these methods also cannot handle geometrical distortion corrections or arbitrary rotations or scale changes. In general it is also not possible to give each input pixel its own weighting.

2.2. Drizzling

During the preparation for the major HST/WFPC2 imaging campaign of the Hubble Deep Field North in late 1995 (Williams et al. 1996) a need was recognised for more flexible and efficient image reconstruction software suitable for the large volumes of undersampled WFPC2 data. The aim was to minimise the resolution degradation during combination, to handle the known geometrical distortion of WFPC2, to guarantee photometric fidelity and to combine and propagate pixel weights optimally. As the data volumes were large and time limited a reasonably high throughput was essential. The resulting IRAF implementation was of a method called variable pixel linear reconstruction but normally called "drizzling" (Fruchter & Hook 1997, 1998).

Drizzling is a "forward" method unlike typical interpolation methods. Each pixel of the input images is "shrunk" by a user-specified amount (known as the *pixfrac*) and the corners of this smaller square are transformed onto the output pixel grid using knowledge of the geometrical distortion and any shift, rotation and scale specified. The overlap of this quadrilateral with the pixels of the output is calculated and the data values are combined using an optimal weighting scheme in which the input weight depends on the weight assigned to the input pixel as well as the size of the overlap with the output pixel under consideration. The method is illustrated in Figure 2.

Drizzling, although conceptually very simple, is fast and flexible. Large images, mosaicing, arbitrary geometrical distortions can be easily handled. Input pixels can be individually weighted and these weights are optimally combined and propagated to a separate output weight image. The noise characteristics of the resultant output images are understood and simple formulae are available which give the ratio of the noise of a drizzled output image to the case of no noise correlation. This information is very valuable if the drizzled output

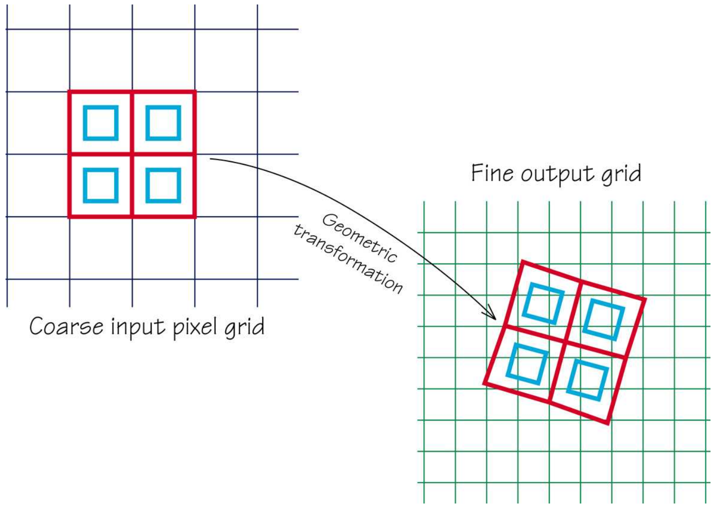

Figure 2. Illustration of how the drizzling method transforms an input pixel onto the selected output grid and showing the pixel shrinkage and general geometric distortion which can be included.

images are to be passed to object detection and classification software such as SExtractor (Bertin & Arnouts 1996) which needs an estimate of the noise. The average FWHM of a point-source in an output drizzled image may be estimated by adding in quadrature the width of the incident optical PSF, the width of the PRF and the *pixfrac* when all three quantities are expressed in the same units. This rule-of-thumb predicts a FWHM of 3.3 pixels for the case of the HST/WFC/F606W HDF-S combined images (which have 0.04" pixels) in close agreement with the measured width of (the rare) stars in this image.

There is a widely used drizzle implementation in IRAF. It has become the standard method for the combination of dithered HST imaging data and has also been used for other data (e.g., ISOCAM, ESO Imaging Survey). An example of the application of drizzling to the deepest optical image of the sky yet taken (Gardner et al. 2000) is shown in Figure 4. This is a combination of 193 unfiltered STIS CCD images with a total exposure time of 155ks. The large number of images resulted in comprehensive sub-pixel sampling and allowed a combination which was equivalent to "interlacing".

The implementation of drizzling using the scheme shown in Figure 2 is just one possibility. Different kernels for distributing weight on the output grid could be used and might have advantages for some applications. Gilliland et al. (1999) have used a method of their own which is similar to classic drizzling but uses a Gaussian kernel. This and other options will be included in a future release of drizzle.

On the other hand drizzling may be criticised in various ways: the choice of the *pixfrac* parameter is somewhat arbitrary; there is a small amount of space-variant smoothing of the output image causing noise-correlations on small scales; the effective interpolation scheme applied is a variant of linear interpolation results in some aliasing. Finally, as with all linear reconstruction methods, drizzling makes no attempt to reduce the loss of resolution resulting from the convolution with either the PSF or the PRF.

The actual combination of the images is only part of the processing of dithered data sets. It is also necessary to measure the shifts between frames accurately as well as detect and flag artifacts so that they do not contribute to the output image. This is particularly difficult when all the data frames have different pointings and it is not possible to detect artifacts using conventional methods. A package of tools for handling dithered HST data is available as the **dither** package in STSDAS (Fruchter et al. 1997). It has also proved possible to use drizzling along with other tools to register such images, detect and flag bad pixels in the input images and then do an optimal combination. Figure 3 gives an example of the application of this technique. Gonzaga et al. (1998) have compiled a "cookbook" where comprehensive worked examples of applying the **dither** package to a variety of realistic datasets are presented.

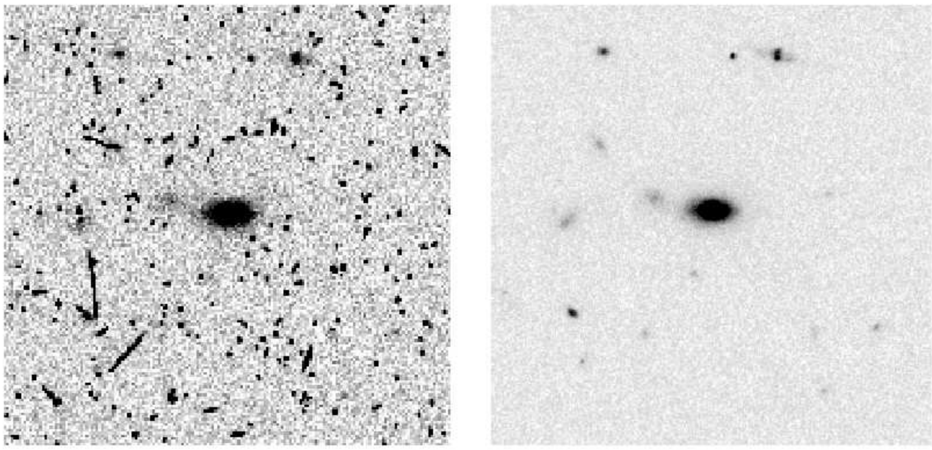

Figure 3. Example of the cosmic-ray rejection and combination of a set of 12 deep WFPC2 images when each has a different pointing. The image on the left is one of the inputs and shows many cosmic-ray hits. The combination on the right, which has a finer pixel grid, is well-cleaned of artifacts and clearly improved small-scale structure.

2.3. Algebraic Fourier Combination

Tod Lauer (1999a) has recently looked at the problem of combining dithered undersampled images in a fresh way with the aim of avoiding some of the problems of the methods discussed so far. The aim was reconstruction of a "super-image" which is Nyquist sampled without the small and space varying blurring which is inevitable in methods such as drizzling. His method works in Fourier space and follows from earlier work on one-dimensional sampled data by Bracewell (1978).

The Fourier transform of an undersampled dataset is periodic with the "satellites" overlapping each other. This overlap is the cause of aliasing and leads to artifacts in data space. However, when multiple dithered input images are available a linear combination of the Fourier transforms may be derived in

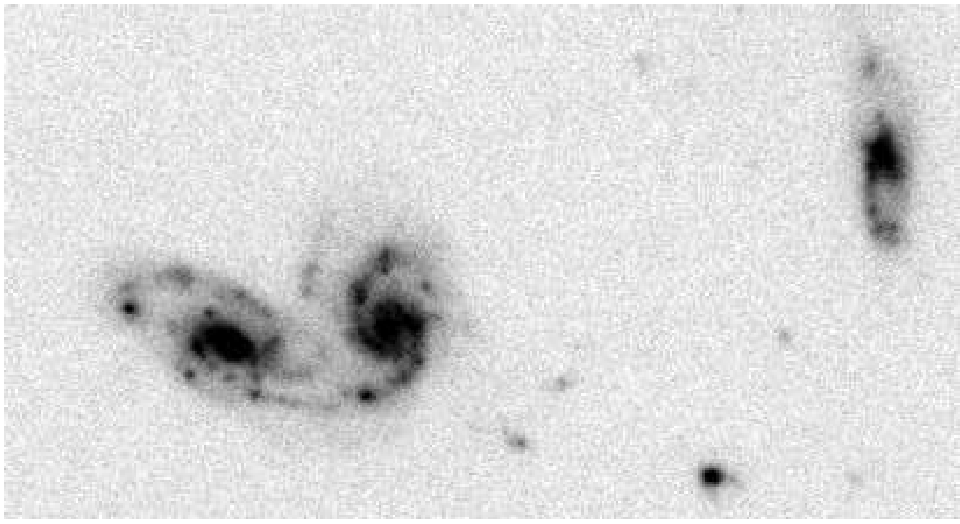

Figure 4. Example of the drizzled combination of many deep HST/STIS CCD unfiltered images of part of the Hubble Deep Field South The two spiral galaxies at the lower-left are separated by less than 2".

which the aliasing is suppressed and the Fourier transform of a critically sampled "super-image" computed.

This method is probably the best currently available for reconstructing fine scale detail of a small region of the sky free of aliasing artifacts in the case of well-dithered data sets. Because the combination is done in Fourier space it is very difficult to include geometric distortion correction and flexible pixel weighting. It is proposed that these steps can be separated from combination itself and done as pre or post-processing. Unfortunately there is no current common-user implementation which limits its current applicability.

2.4. Non-linear Iterative Methods

All the methods described so far attempt to reconstruct the convolution of the sky intensity distribution with both the PSF and PRF and suppress the effects of undersampling. A more ambitious goal is to, in addition, try to remove some of the blurring by the applications of image restoration techniques. One simple way in which this may be done is to use a multiple input-channel generalisation of the Richardson Lucy maximum likelihood algorithm (Adorf & Hook 1995; Hook & Adorf 1995). In this case the shifts may be handled by assigning a shifted PSF to each input image. Rotation and more general geometric distortions cannot be included and there is an additional disadvantage that the resulting images have strongly correlated noise characteristics and may suffer from artifacts if too much super-resolution is attempted. Figure 5 shows a comparison of this method (in the ACOADD IRAF implementation) along with standard drizzled combinations.

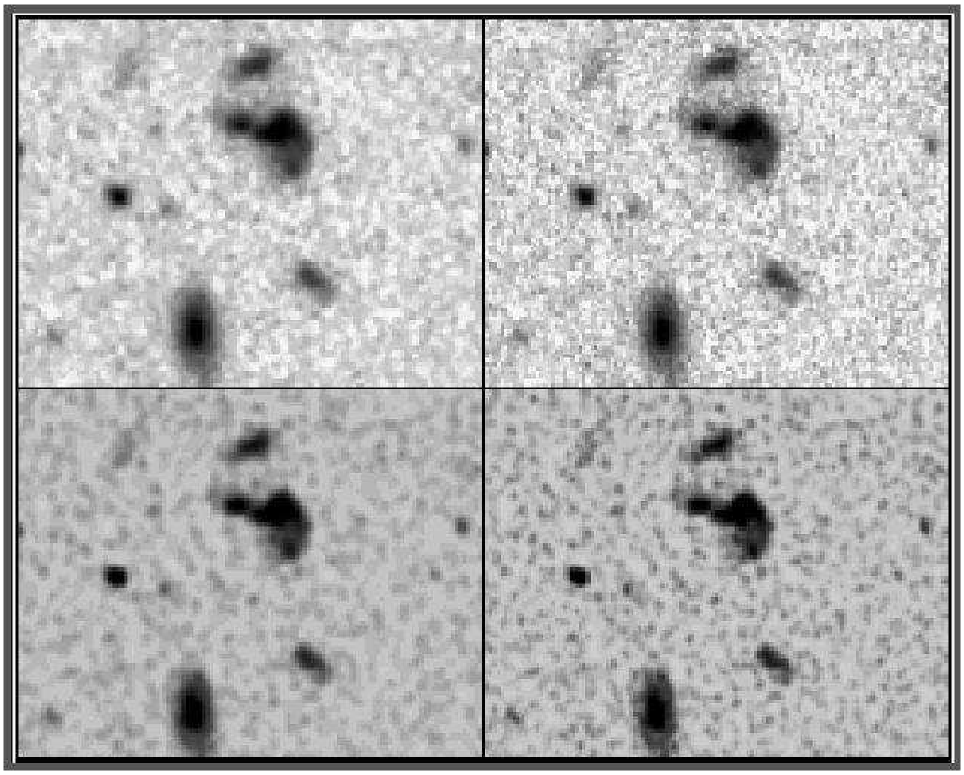

Figure 5. A comparison of combinations of deep HST NICMOS Camera 3 data. There were 9 evenly-spaced dither positions. The upper combinations were done with Drizzle and the lower ones with ACOADD. The upper left had $pixfrac = 1.0$ and the upper right $pixfrac = 0.5$. The lower left was smoothed with a Gaussian of $\sigma = 1.0$ output pixels and that at the lower right with $\sigma = 0.7$. Data courtesy Mark Dickinson (STScI).

3. The Revenge of the Pixel-Response Function

Each individual pixel of any detector will have variations of sensitivity across its surface. When the incident PSF is well sampled such variations within the PRF have only very small effects on photometry and astrometry and can normally be neglected. Unfortunately this is no longer true in the undersampled images cases being considered here and very significant photometric differences may occur between the case of a point-source being imaged on the centre or the edge of a pixel.

This effect can be mapped directly if many dithered exposures of the same field are available. An example of this is given in Figure 6 for the case of the NICMOS observations of the Hubble Deep Field South. Alternatively the effect of the PRF on photometry may be deduced by reconstructing the convolution of the PSF and PRF using Lauer's method and then moving a sampling grid of δ-functions and summing up to obtain a two-dimensional map of the measured response of a point-source at different sub-pixel positions. Applications of these methods to the WFPC2 and Camera 3 of NICMOS are described in Storrs et al. (1999) and Lauer (1999b) where more details and suggested schemes for correcting this effect are presented.

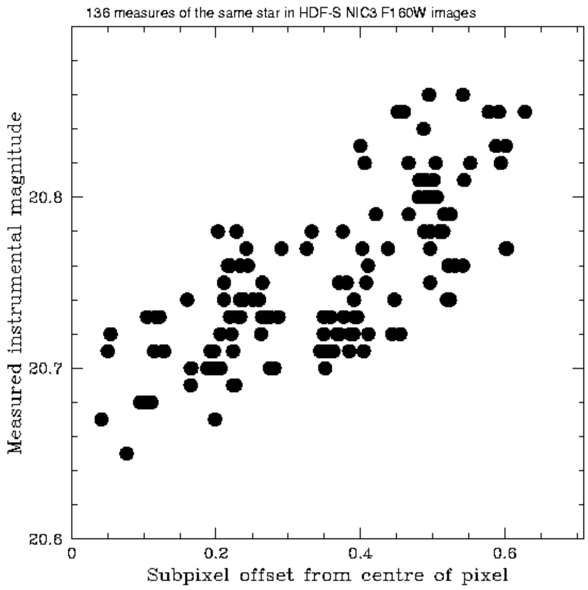

Figure 6. Variation of measured brightness of a single star image in 136 separate dithered images of the HDF-S NICMOS field using Camera 3 plotted as a function of sub-pixel distance of the centre of the star from the centre of the pixel.

4. Conclusions & Future Developments

The advent of large volumes of high-quality dithered undersampled images from HST and elsewhere has prompted the development of new tools for optimal reconstruction of such data sets. These methods have been effective but there is much room for further enhancements. The optimal choice of method will be dictated by the scientific problem being tackled. The next few years will see the installation of the Advanced Camera for Surveys on HST which will produce a much larger volume of data than the current cameras and suffer from extreme geometrical distortion and consequent pixel shape and size variations around the field. Looking further ahead the NGST will continue this process and bring new problems of its own.

References

Adorf, H.-M., & Hook, R. 1995, "High resolution images from multiple 'dithered' frames", in: Proc. ST-ECF workshop on "Calibrating and understanding HST and ESO instruments", (Garching: European Southern Observatory), 251

Bertin, E. & Arnouts, S. 1996, A&AS, 117, 393

Bracewell, R. N. 1978, "The Fourier Transform and Its Applications" (New York: McGraw-Hill)

Devillard, N., 1999, "Infrared Jitter Imaging Data Reduction: Algorithms and Implementation", in ASP Conf. Ser., Vol. 172, Astronomical Data Analysis Software and Systems VIII, ed. David M. Mehringer, Raymond L. Plante, & Douglas A. Roberts (San Francisco: ASP), 333

Fruchter, A. S., & Hook, R. N. 1997, "A novel image reconstruction method applied to deep Hubble Space Telescope Images", Invited paper, in Applications of Digital Image Processing XX, ed. A. Tescher, Proc. S.P.I.E. vol. 3164, 120

Fruchter, A. S., Hook, R. N., Busko, I. C., & Mutchler, M. 1997,"A Package for the Reduction of Dithered Undersampled Images", in proceedings of the 1997 HST Calibration Workshop, eds. Stefano Casertano, Robert Jedrzejewski, Charles D. Keyes, and Mark Stevens, (Baltimore: STScI), 518

Fruchter, A. S., & Hook, R. N. 1998, "A Method for the Linear Reconstruction of Undersampled Images", astro-ph/9808087, submitted to PASP

Gardner J. P., Baum, S. A., Brown T. M., Carollo, C. M., Christensen, J., Dashevsky I., Dickinson M. E., Espey B. R., Ferguson H. C., Fruchter A. S., Gonnella A. M., Gonzalez-Lopezlira R. A., Hook R. N., Kaiser M. E., Martin C. L., Sahu K. C., Savaglio, S., Smith T. E., Teplitz H. I., Williams R. E., & Wilson, J. 2000, "The Hubble Deep Field South - STIS Imaging", AJ, in press

Gilliland, R. L., Nugent, P. E., & Phillips, M. M. 1999, "High-Redshift Supernovae in the Hubble Deep Field", ApJ, 521, 30

Gonzaga, S. et al. 1998, "The Drizzling Cookbook", STScI Instrument Science Report WFPC2 98-04

Hook, R. N., & Adorf, H.-M. 1995, "Methods for combining 'dithered' WFPC-2 images", in: Proc. Calibrating Hubble Space Telescope: Post Servicing Mission, Baltimore, MD, Space Telescope Science Institute, 341

Lauer, T. R. 1999a, "Combining Undersampled Dithered Images", PASP, 111, 227

Lauer, T. R. 1999b, "The Photometry of Undersampled Point Spread Functions", PASP, 111, 1434

Storrs, A., Hook, R., Stiavelli, M., Hanley, C. & Freudling, W. 1999, "Camera 3 Intrapixel Sensitivity", STScI Instrument Science Report NICMOS-99-005

Williams, R. E., Blacker, B., Dickinson, M., Dixon, W. V., Ferguson, H. C., Fruchter, A. S., Giavalisco, M., Gilliland, R., Heyer, I., Lucas, R. A., McElroy, D. B., Petro, L., Postman, M., Adorf, H-.M., & Hook, R. N. 1996, "The Hubble Deep Field: Observations, Data Reduction, and Galaxy Photometry", AJ,112, 1335.