

Linear Reconstruction of the Hubble Deep Field

A.S. Fruchter¹ and R.N. Hook²

¹ Space Telescope Science Institute

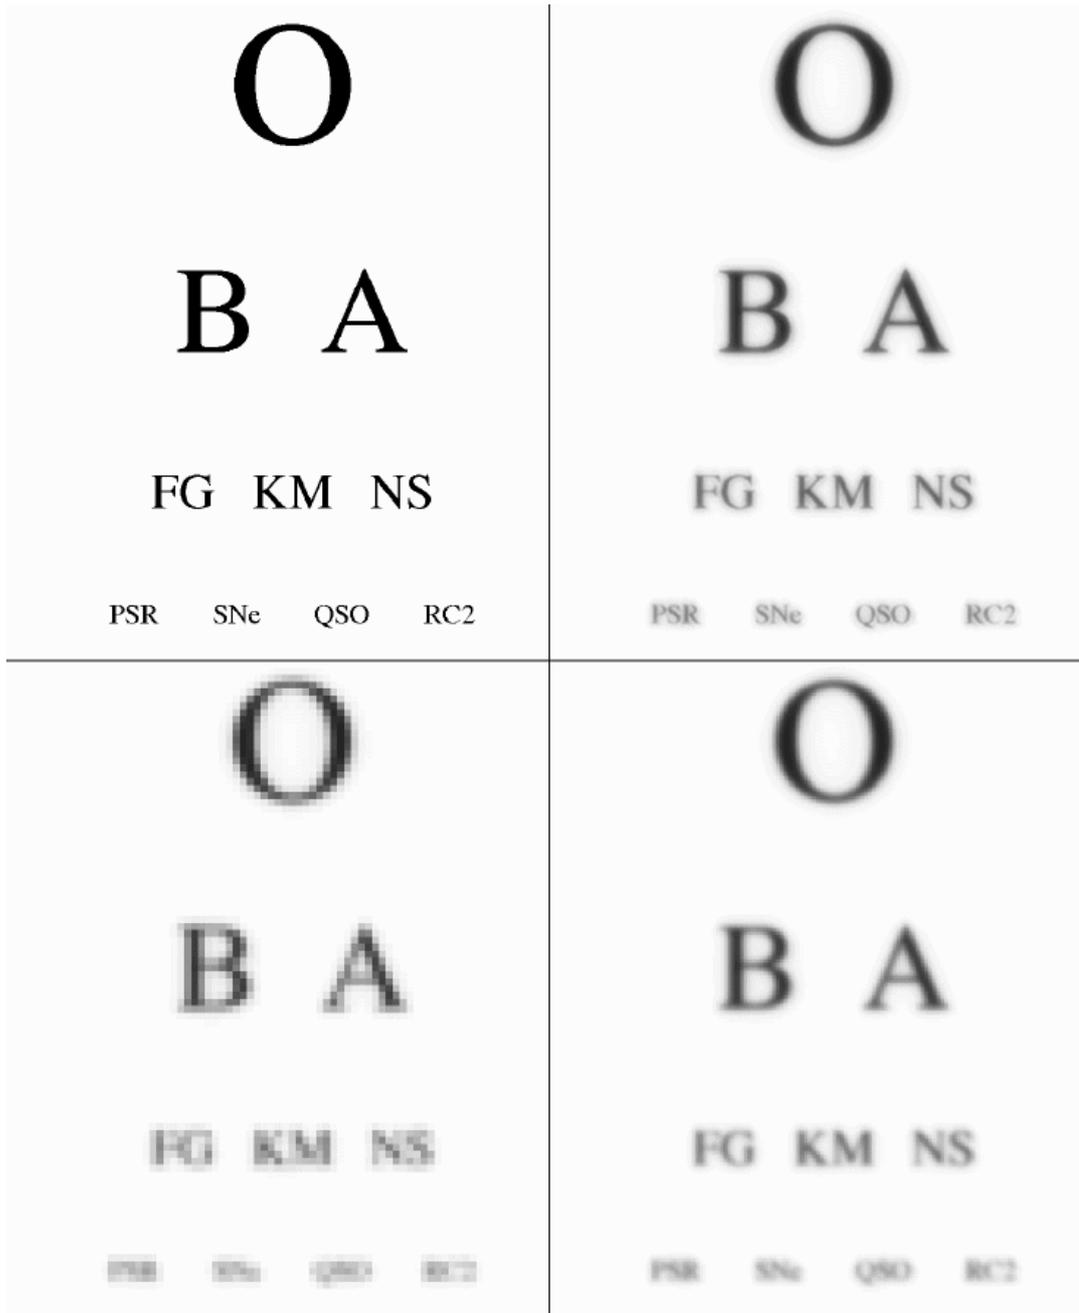
² Space Telescope-European Coordinating Facility

Abstract

We have developed a method for the linear reconstruction of an image from undersampled, dithered data, which has been used to create the distributed, combined Hubble Deep Field images. The algorithm, known as Variable-Pixel Linear Reconstruction (or informally as "drizzling"), preserves photometry and resolution, can weight input images according to the statistical significance of each pixel, and removes the effects of geometric distortion both on image shape and photometry. In this paper, the algorithm and its implementation are described, and measurements of the photometric accuracy and image fidelity are presented. In addition, we describe early attempts to use drizzling to combine dithered images in the presence of cosmic rays.

Introduction

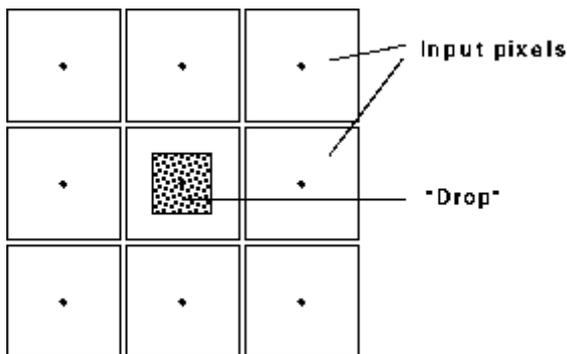
Although the optics of WFPC2 now provide a superb PSF, the detectors at the focal plane undersample the image. This problem is most severe on the three WF chips, where the width of a pixel equals the FWHM of the optics in the near-infrared, and greatly exceeds it in the blue. The effect of undersampling is illustrated by the "Great Eye Chart in the Sky" below. In the upper left hand corner one sees the "true" image, as it would be seen by a telescope of infinite aperture. In the upper right, the image has been convolved with the red PSF of the WF2 and in the lower left it has been subsequently sampled by the WF CCD. The loss of spatial information is immediately obvious.



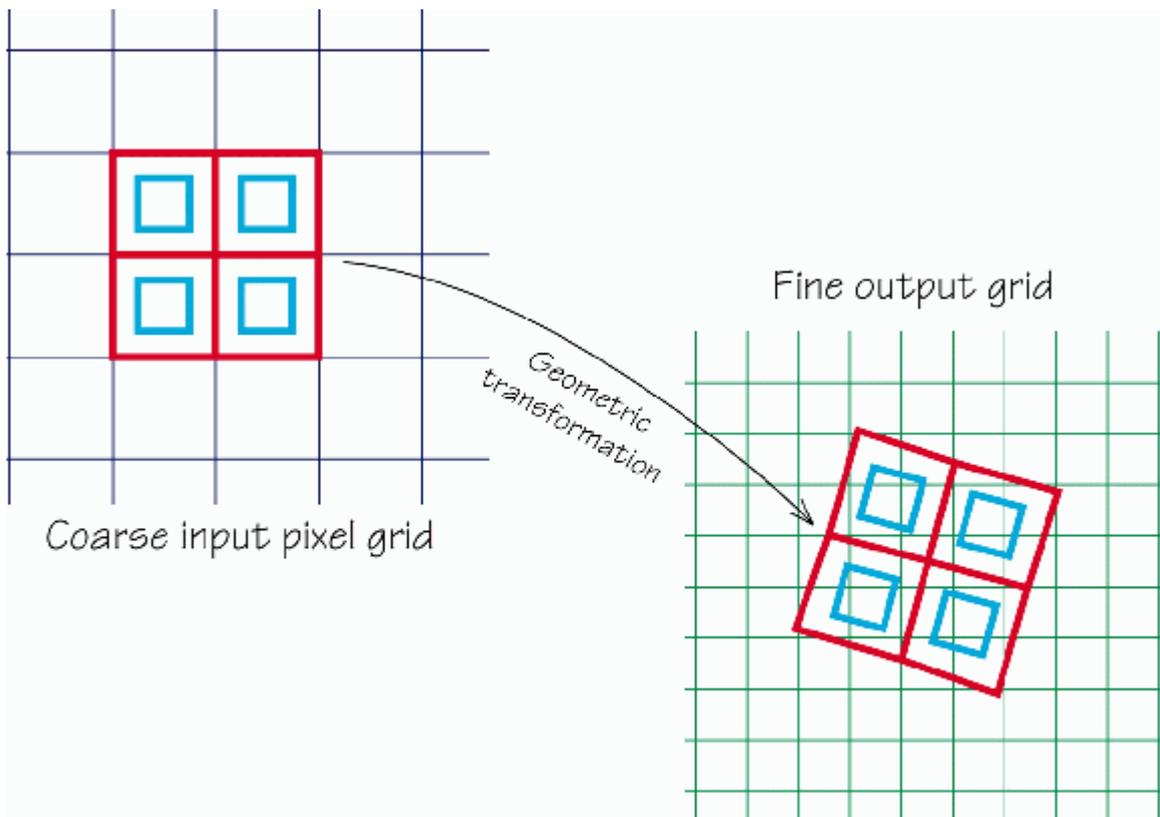
Fortunately, much of the information lost in sampling can be recovered. In the lower right we display the image recovered using one of the family of techniques we refer to as "linear reconstruction." The most commonly used of these techniques are shift-and-add and interlacing. The image in the lower right corner has been restored by interlacing a 3x3 array of dithered images. However, due to poor placement of the sampling grid or the effects of geometric distortion, true interlacing of images is often infeasible. On the other hand, the other standard technique, shift-and-add, convolves the image yet again with the original pixel, adding to the blurring of the image and the correlation of the noise. (The deterioration in image quality between the upper and lower right hand images is due entirely to convolution with the WF pixel). Here we present a method which has the versatility of shift-and-add yet largely maintains the resolution and independent noise statistics of interlacing.

The Method

The **drizzle** algorithm is conceptually simple. Pixels in the original input images are mapped into pixels in the subsampled output image, taking into account shifts and rotations between images and the optical distortion of the camera. However, in order to avoid convolving the image with the large pixel "footprint" of the camera, we allow the user to shrink the pixel before it is averaged into the output image, as shown in the figure below.

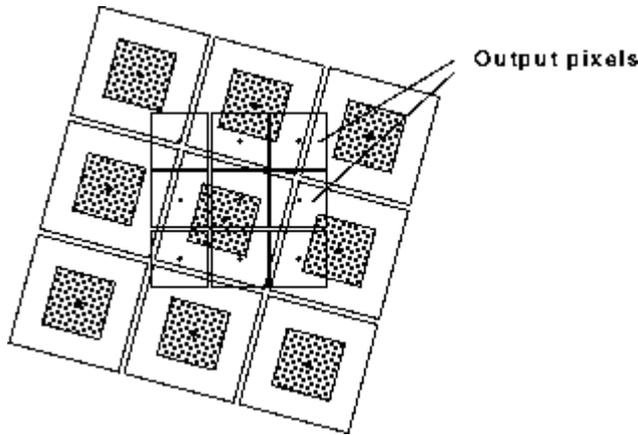


The new shrunken pixels, or "drops", rain down upon the subsampled output image.



In the case of the HDF, the drops had linear dimensions one-half that of the input pixel -- slightly larger than the dimensions of the output subsampled pixels. The flux in each drop is divided up among the overlapping output pixels in proportion to the areas of overlap. Note that if the drop size is sufficiently small not all output pixels have data added to them from each input image. One must therefore choose a drop size that is

small enough to avoid degrading the image, but large enough that the after all images are "drizzled" the coverage is fairly uniform.



In averaging the input image with the output image, the size of the drop is further adjusted to reflect the geometric distortion of the camera before the overlap of the drop with pixels in the output image is determined. When a drop with value i_{xy} and user defined weight w_{xy} is added to an image with pixel value I_{xy} , weight W_{xy} , and fractional pixel overlap $0 < a_{xy} < 1$, the resulting value of the image I'_{xy} and weight W'_{xy} is

$$W'_{xy} = a_{xy}w_{xy} + W_{xy}$$

$$I'_{xy} = \frac{a_{xy}i_{xy}w_{xy} + I_{xy}W_{xy}}{W'_{xy}}$$

Drizzling and Photometric Accuracy

The field flattener of the WFPC2 geometrically distorts the images: pixels at the corner of each CCD subtend less area on the sky than those near the centers of the CCD. However, after application of the flat field, a source of uniform surface brightness on the sky produces uniform counts across the CCD. Therefore point sources near the corners of the chip are artificially brightened compared to those in the center of the chips.

In order to test the effect of drizzling on stellar photometry, we created a 19x19 grid of artificial stars on a 4 times oversampled grid using the Tiny Tim WFPC PSF modelling code and then convolved these stars with a narrow Gaussian which approximates the smearing caused by the cross-talk between neighboring pixels in the WFPC. These images were then multiplied by the Jacobian of the geometric distortion of the WF3 camera, to represent the effect of the geometric distortion on point source photometry. This image was then shifted and sampled on a 2x2 grid and the results combined using the **drizzle** algorithm with an output pixel size one-half that of the original WF pixels, and a drop size with linear dimensions 0.65 that of the WF pixel. The geometric distortion of the chip was removed during drizzling using the polynomial model of

Trauger *et al.* The amount of data and dithering pattern, therefore, resemble ones that a typical user might produce. (In contrast, the HDF contained 11 different pointings.)

To determine the effectiveness of drizzling in correcting the photometric effects of geometric distortion, we then obtained aperture photometry on the stars in one of the original input images and on the stars in the output drizzled image. In the image to the lower left we display the results obtained from the input image. The photometric measurements of the 19x19 stars are represented by a 19x19 pixel image. The effect of the distortion on the photometry is obvious -- the stars in the corners are up to ~4% brighter than those in the center of the chip. The image on the right similarly displays the results of aperture photometry on the 19x19 grid of stars after drizzling. The effect of geometric distortion on the photometry is dramatically reduced: the r.m.s. photometric variation in the drizzled image is 0.004 mags.

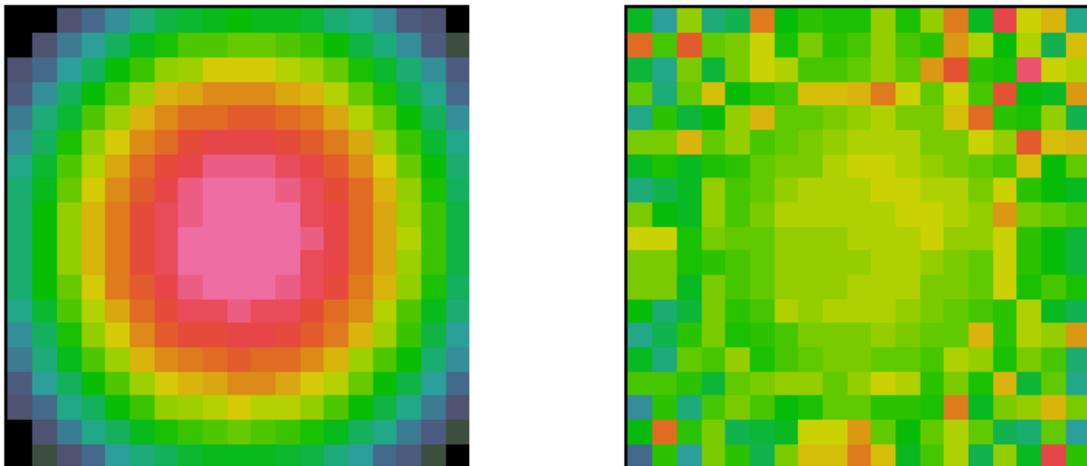


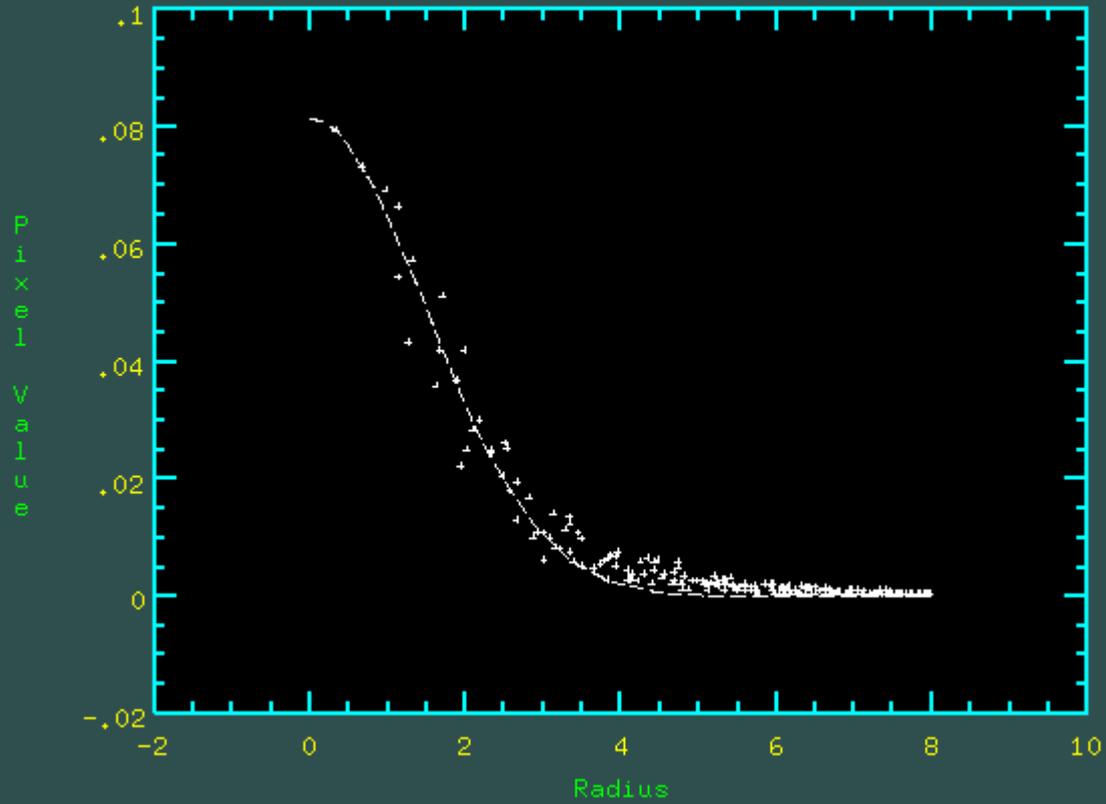
Image Fidelity

The **drizzle** algorithm was designed to obtain optimal signal-to-noise on faint objects while preserving image resolution. To a some degree, however, these goals are incompatible. Indeed, image restoration procedures which attempt to remove the blurring of the PSF and the pixel by enhancing the high frequencies in the image (such as such as the Richardson-Lucy and MEM methods) directly exchange signal-to-noise for resolution. In the **drizzle** algorithm we have made no compromises on signal-to-noise. In particular, the weight of an input pixel in the final output image is entirely independent of its position on the chip. Therefore, if the dithered images do not uniformly sample the field, the "center of light" in an output pixel may be offset from the center of the pixel, and, if the sampling is not uniform, the offset may vary between adjacent pixels. This effect is seen in the HDF images, where some pointings were not at the requested position or orientation. Furthermore, the very large dithering offsets used in the HDF, combined with geometric distortion, produce a sampling pattern that varies across the field.

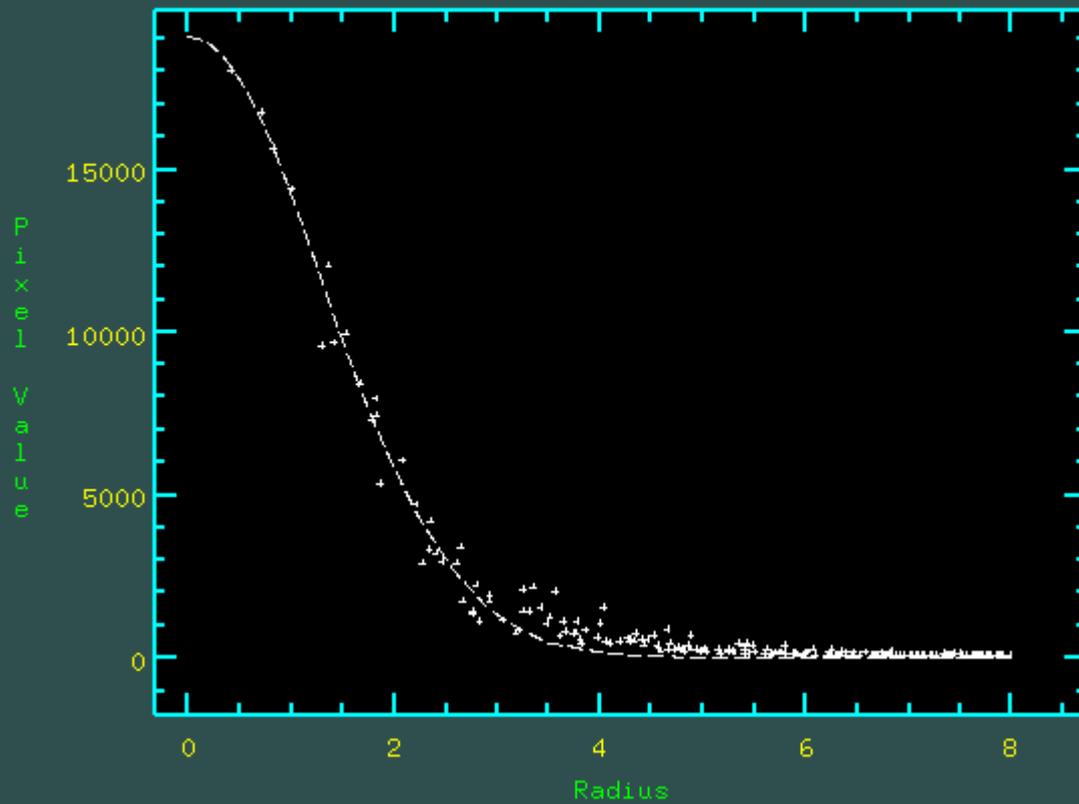
Below are shown two PSFs. The upper PSF is taken directly from the HDF F450W drizzled image. The PSF shows substantial variations about the best fit Gaussian due to the effects of non-uniform sampling (note, however, that these variations do not noticeably affect aperture photometry performed with a radius greater than or equal to 5

output pixels -- that is 2 WF pixels). The lower PSF is bright star taken from a deep image with a nearly perfect four-point dither. The uniform sampling produces a smooth PSF. The difference in the apparent widths of the PSFs is due to the use of larger output pixels in the second image than in the HDF (0."05 vs. 0."04).

NOAO/IRAF V2.10.4EXPORT fruchter@electra.stsci.edu Wed 17:38:58 28-Feb-96
F450_g2_dr: Radial profile at 235.92 1496.67
F450_G2_DR[1/1]



NOAO/IRAF V2.10.4EXPORT fruchter@electra.stsci.edu Tue 17:30:49 18-Feb-97
sum_noc: Radial profile at 1003.66 274.25
SUM_NOC[1/1]



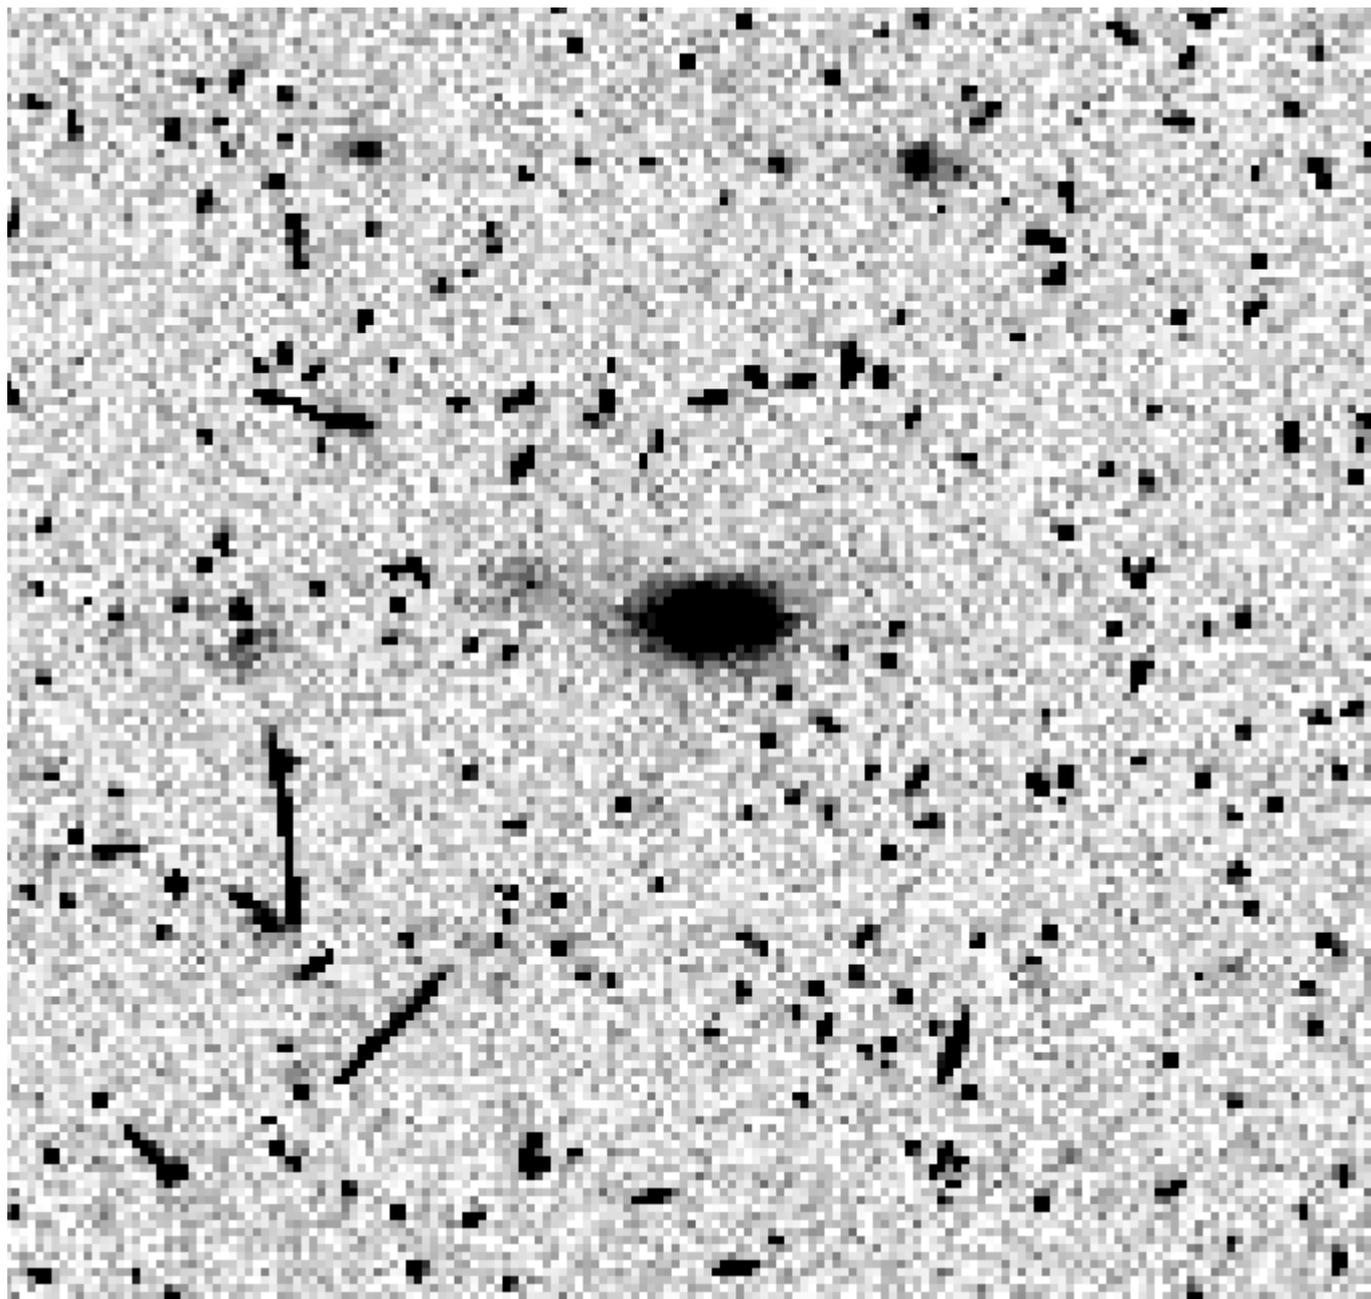
Cosmic Ray Removal

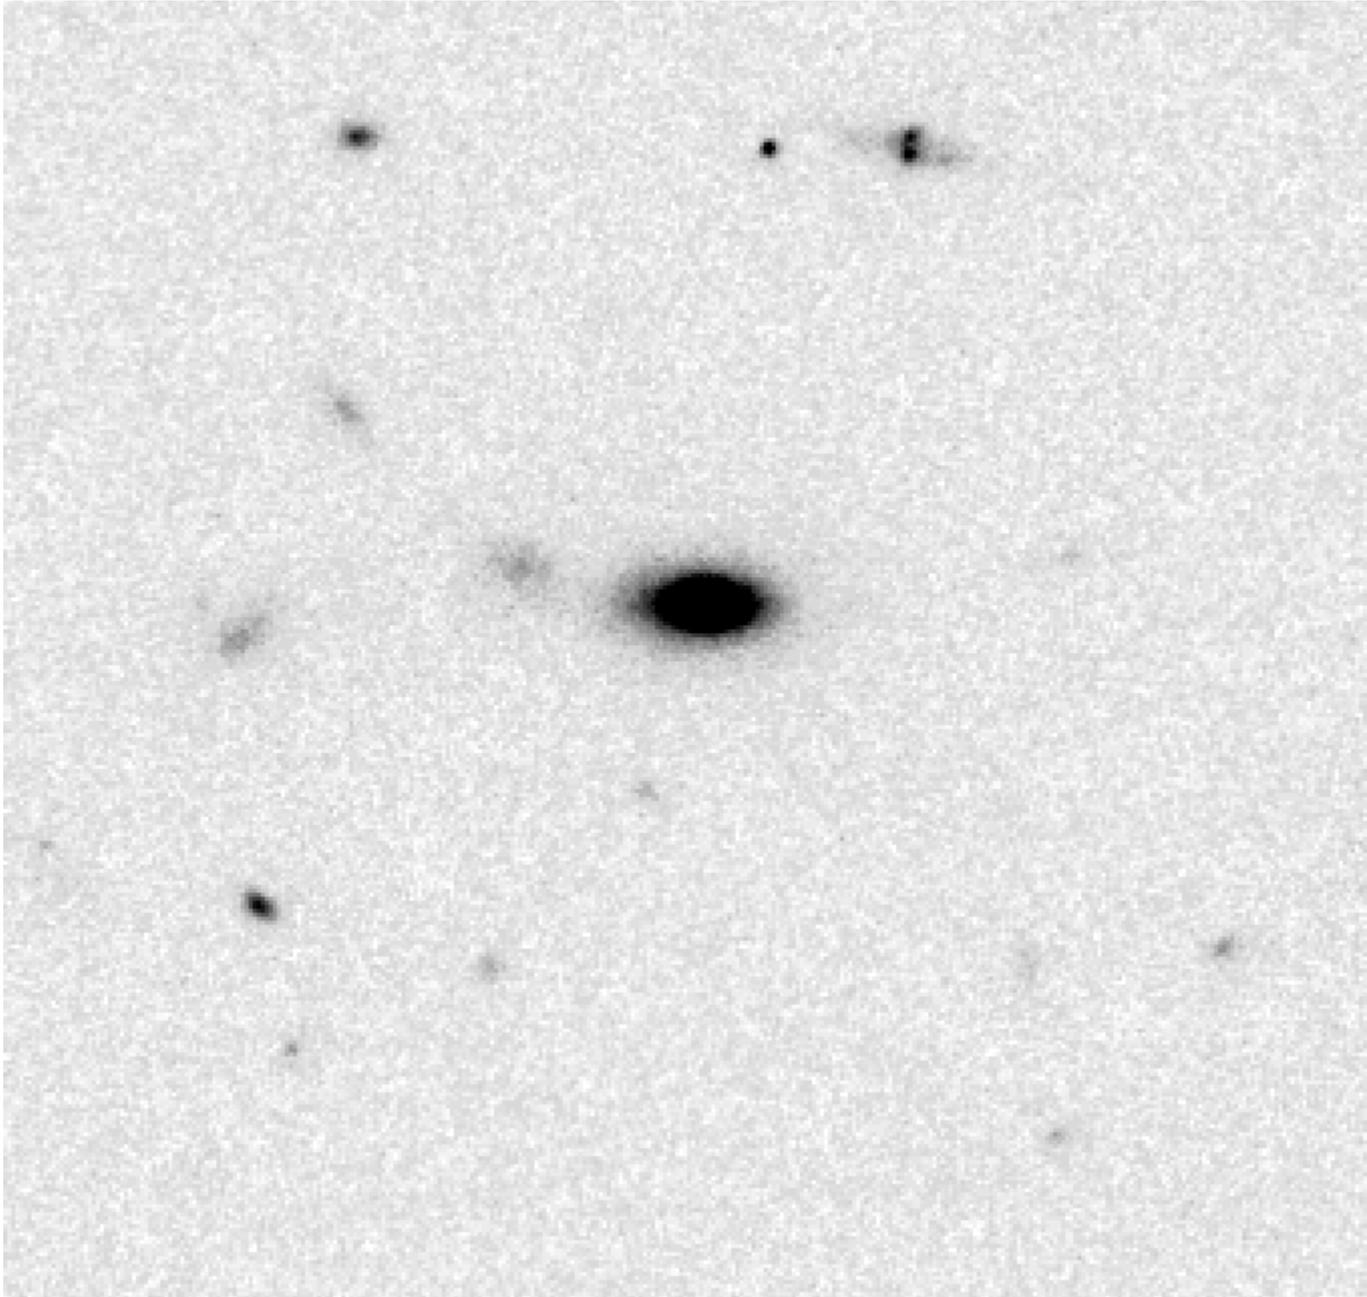
Few HST observing proposals have sufficient time to take a number of exposures at each of several dither positions. Therefore, if dithering is to be of wide-spread use, one must be able to remove cosmic rays from data where few, if any, images are taken at the same position on the sky. We have therefore been examining the question of whether we can adapt the drizzling procedure to the removal of cosmic rays.

Although the removal of cosmic rays using **drizzle** is still very much a work in progress, we have developed the following procedure which appears quite promising.

- Drizzle each image onto a separate sub-sampled output image, but preserve the initial pixel size (thus, for instance, an initial pixel might cover the area four output pixels, but because of the effects of the shifting and geometric distortion, the output image contains sub-pixels which are a weighted average of adjacent pixels, rather than just a block replication of the input).
- Take the median of the output drizzled images.
- Map the median image back to the input plane of each of the individual images, including the image shifts and geometric distortion. This is done using a program named **blot** (blotting is, in effect, the inverse of drizzling).
- Take the spatial derivative of each of the blotted output images.
- Compare each original image with the corresponding blotted image. Where the difference is larger than can be explained by noise statistics, or the flattening effect of taking the median or perhaps an error in the shift (the magnitudes of the latter two effects are estimated using the image derivative), the suspect pixel is masked.
- Repeat the previous step on pixels adjacent to pixels already masked, using a more stringent comparison.
- Finally, drizzle the input images onto a single output image using the pixel masks created in the previous steps.

Below we display two figures. The upper figure is part of a single 2400 second exposure taken with the F814W filter and the WF2 CCD. The lower image shows the results of applying the above procedure to twelve such exposures, nearly all of which had been dithered from the central position by approximately 50 pixels. No two of the exposures were taken at the same position on the sky. The gain in effective resolution produced by dithering is evident in the "double" nucleus of the galaxy in the upper right. The two nuclei are separated by 0.2 arcseconds.





Most of the information displayed on this page can be found in a paper, written by Richard and myself, for an SPIE conference. This version is easily printable, and can be retrieved from [astro-ph](#), an "e-print" archive of astronomical papers.