The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Band-limited Imaging with Undersampled Detectors

Andrew S. Fruchter

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. Over the past decade "Drizzle" has become a de facto standard for the combination of HST images. However, the drizzle algorithm was developed with small, faint, partially-resolved sources in mind, and is not the best possible algorithm for high signal-to-noise unresolved objects. Here, a new method for creating band-limited images from undersampled data is presented. The method uses a drizzled image as a first order approximation and then rapidly converges toward a band-limited image which fits the data given the statistical weighting provided by the drizzled image. The method, named iDrizzle, for iterative Drizzle, eliminates the small high-frequency artifacts that can be introduced by drizzling. The method works well in the presence of geometric distortion, and can easily handle cosmic rays, bad pixels, or other missing data. It can combine images taken with random dithers, though the number of dithers required to obtain a good final image depends in part on the quality of the dither placements.

1. Introduction

The competition between the desire to have a wide field-of-view and the limitations on the number of available pixels – caused by a desire to reduce cost and complexity or to lessen read-noise – means that astronomical detectors are often undersampled. In order to fully sample an image with an undersampled detector one must dither and combine multiple images. However, distortions in the field-of-view may make it impossible to perform shifts that equally well sample different parts of the detector. In practice then the combined pixels from dithering of astronomical detectors often produce irregular sampling of the image plane.

In order to combine the irregularly sampled data from the Hubble Deep Field *HDF* (Williams et al. 1996), a new image algorithm, Drizzle (Fruchter and Hook 2002) was developed. Drizzle combines dithered images in a statistically optimal fashion. However, as can be seen in Figure 1, Drizzle adds small high-frequency artifacts to the image. On scales larger than an original pixel, these rapidly average out. Thus for the prime purpose of the *HDF*, the study of small faint galaxies, Drizzle is an excellent algorithm. However, for the analysis of high signal-to-noise images of point sources, or other cases where preservation of the true point spread function (PSF) was essential, one might prefer an algorithm which more exactly reproduced the true band-limited image that is formed in the detector. While there is no known method to analytically reconstruct an irregularly-sampled band-limited image, here we present an iterative method which performs. In essence it is an iterative approximation to the band-limited image, using Drizzle to ensure that the data are combined using the full statistical power of the data. In the absence of noise this method converges directly to the true image. In the presence of noise a small increase in the statistical noise is produced, but the systematic high-frequency noise introduced by Drizzle is removed.

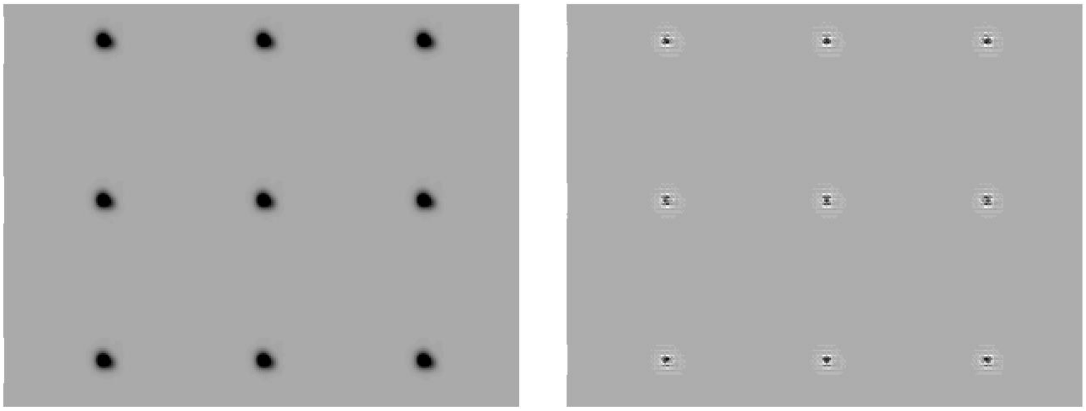

Figure 1: On the left a series of synthetic ACS PSFs. At the right, a drizzled approximation of the image subtracted from the original. The largest residuals are slightly greater than 10% of peak. Drizzle smoothes the PSF and adds high frequency noise. The smoothing is reproducible; thus if a PSF slightly larger than the original is acceptable, the smoothing is not a significant issue. Similarly, if one is measuring properties of the image on scales larger than a couple of original pixels, the high frequency noise largely averages out. Thus Drizzle is well suited for aperture photometry with apertures larger than 2 original pixels, or galaxy photometry. If one wished to reconstruct a true instrumental PSF, or if one wishes to do photometry through PSF fitting, another method might be preferred.

2. The Method

The spatial frequencies in an image are limited by the optics of the telescope to be $\leq D/\lambda_s$, where D is the (maximum) aperture of the telescope and λ_s is the shortest wavelength in the passband. Thus the spatial frequencies in an images are band-limited. One might imagine then that one could take the total set of irregularly sampled data which everywhere locally meets the Nyquist criterion and do a Fourier transform, remove any frequencies above the cutoff frequency, and Fourier transform back to get the true band-limited image. Unfortunately, a direct Fourier transform of an irregularly-sampled data set throws a great deal of power out of the original passband, and thus this method fails terribly. Surprisingly, it is more effective to first put the data onto a regular Nyquist grid by simply taking the value of the nearest neighbor before doing the Fourier transform. The inverted, band-limited function turns out to be a much truer approximation than the direct transform case. This approximation is know as the Voronoi approximation, and more information on it, and the ideas discussed in this paragraph can be found in the tutorial by Tobias Werther (2006).

Now the Voronoi approximation is a band-limited function and thus can be sinc interpolated to the irregular grid of the data. One can therefore subtract the Voronoi approximation from the original function at all of the data points. Furthermore, this smaller difference function is itself a band-limited function, so one can repeat the process and get a further refined approximation to the underlying band-limited function. This procedure is known to converge geometrically (Werther 2006).

The nearest neighbor approximation is, however, far from ideal for astronomical imaging. Only one data sample is used at any point on the regular grid, even if several nearby data samples could provide information, and there is no means to weight the data according to its statistical significance. Therefore in the proposed method the nearest neighbor approximation is replaced by Drizzle in each iteration. While Drizzle introduces small artifacts (as does the nearest neighbor) the iterative comparison with the original data serves to remove these, as can be seen in Figure 2.

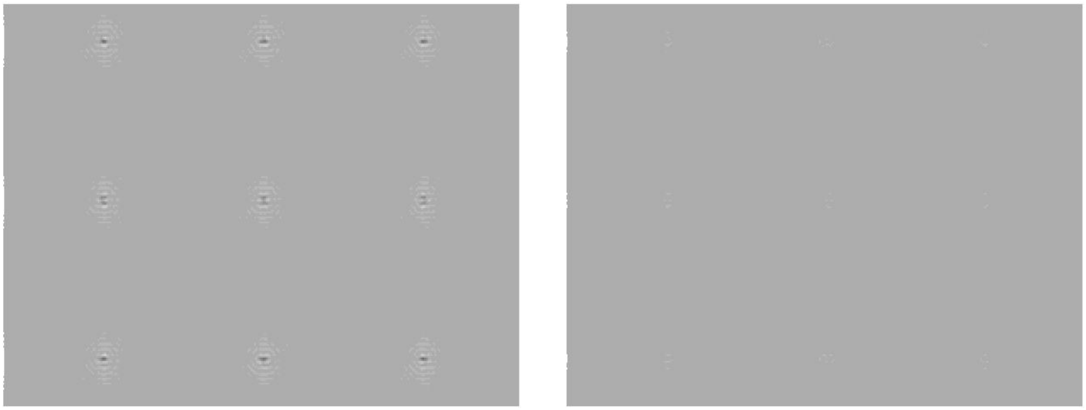

Figure 2: On the left, the residual difference between the true ACS PSF (see Figure 1) and the iterative approximation after the first iteration. On the right the same difference after several more iterations. The rapid reduction of the small artifacts of the Drizzle algorithm is evident. Twelve simulated noiseless ACS images with random pointings were combined.

It is noise which requires the use of Drizzle, rather than the nearest neighbor approximation, but in Figure 2, there is no noise. On the left-hand-side of Figure 3 we therefore show the same subtraction performed between noiseless ACS PSFs and a combination of twelve simulated ACS images of the stellar field, with each star near saturation in a 1200s exposure, with appropriate Poisson and read noise added. This tests the method in a situation of extremely high signal-to-noise. The residuals are close to that expected from noise statistics, and the peak errors are reduced from the Drizzled subtraction by a factor of ~ 20. However, the introduction of noise greatly slows convergence – twenty-four iterations were used to produced the output shown here.

On the right-hand-side of Figure 3 is a central region of the Hubble Ultra-Deep Field (Beckwith et al. 2006). The bright star in this image is near saturation in each of the twelve 1200s individual exposures combined with the new method to form the final image. When this image is mapped back onto the individual input exposures and subtracted, residuals of approximately 2% peak are found under the stellar image. These small but measurable errors are most likely caused by temporal variations in the PSF caused by variation in the insolation of the telescope as it orbits the earth.

3. Noise

3.1. Noise Amplification

Drizzle places the output image exactly where it was observed. But the average weight of an output pixel will not necessarily fall at the center of that pixel, and thus there is a jitter between the represented and effective position of a pixel. Furthermore the peak of a drizzled PSF will never be greater than the greatest value in the appropriate region of the input images. By contrast iDrizzle attempts to predict the true value of the image at the center of the output pixel, and thus the peak of a PSF will often be brighter than any value of the input images in the appropriate region. iDrizzle essentially uses estimates of the derivatives of the data to extrapolate to a position (the center of the output pixel) which is not necessarily exactly sampled in any of the input images. This will produce some noise amplification, which will vary with the quality of the dithering. In typical tests performed on this method the noise amplification has been in the region of 10%.

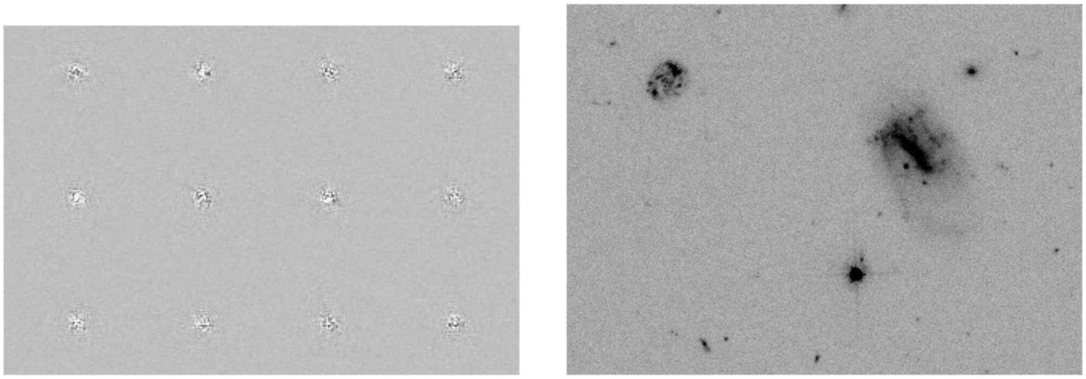

Figure 3: On the left, the same subtraction of the combination of twelve simulated ACS images from the true PSF as seen in Figure 2 but with Poisson and read noise incorporated. On the right, the combination of twelve images from the Hubble Ultra-Deep Field using the new method.

3.2. Correlated Noise

Images produced by Drizzle show correlated noise. Part of this correlation is caused by the drizzling process itself – a non-zero value of `pixfrac` causes Drizzle to place a given input pixel value down on a region of size $p \times s$, where p is the value of `pixfrac` and s is the size of an input pixel. As the iteration proceeds, iDrizzle effectively forces p to zero. But this does not entirely remove correlated noise.

To see why correlated noise remains consider the following situation: the center of an input pixel falls directly on the boundary between two output pixels. Because the image is band-limited (and presumably the output image uses a sampling at least as fine as Nyquist) this input value must affect the predicted value of the two pixels on either side of it. Thus any noise in this input pixel will affect the noise in these two pixels, and thus the noise in these pixels will be correlated. Thus this method reduces but does not eliminate correlated noise.

4. Final Comments

In this short paper, a new method for the combination of dithered astronomical images has been introduced. The method has the ability to handle shifts, distortions and missing data, and converges rapidly to an accurate representation of the underlying image. It requires, however, that the combined images Nyquist sample the image. In the case of random placement of images this may take a minimum of order eight and twelve images for the ACS and WFPC2, respectively, on *HST*.

This article is intended as an outline for iDrizzle. Details – such as apodization of the image and the cut in the Fourier plane – are defered to later and longer presentations. Here, the essence of the method, and its benefits and limitations are presented.

References

Beckwith, S. V. W. et al. 2006, AJ, submitted

Fruchter, A. S. & Hook, R. N. 2002, PASP, 114, 144

Werther, T. 2006, http://www.math.ucdavis.edu/s̃trohmer/research/sampling/irsampl.html

Williams, R. E. et al. 1996, AJ, 112, 1335