# Stacking Images with Drizzle
by Craig Stark

One of the great things that increasing the aperture of a telescope gives us is the ability to resolve finer and finer details in our images. What happens, however, when our telescope optics can resolve fine details (such as the separation between two stars) that are too small to resolve given the size of the pixels on our CCD cameras? For example, what if we imaged a close binary pair whose stars are centered on adjacent or nearly adjacent pixels? An image of this would show several bright, connected pixels, looking somewhat like an oblong star. How might we resolve it into two stars?

One option would be to magnify the image (with something like a Barlow or a longer focal length objective) before the light hits our CCD camera, thus increasing the physical separation of the stars on the CCD array. A second option would be to replace the camera and use one with smaller pixels. What was once a half-pixel separation between the stars might be several pixels if a CCD chip with much smaller pixels were used.

What if we can't change the camera or the telescope? What if we're stuck with the gear and the images we have? We might dream of the next great telescope or the next great camera to help us resolve finer details from our back yard and stare at that credit card while surf the web, browsing ever more powerful (and more expensive!) gear. But what if the telescope and camera were in space, making the cost of changing many orders of magnitude more expensive than the limit on anyone but Bill Gates' credit card?

Undersampling and "Dither"

Drizzle is a technique developed by Andy Fruchter (Space Telescope Science Institute) and Richard Hook (Space Telescope European Coordinating Facility) to solve this problem and to make the most of images from
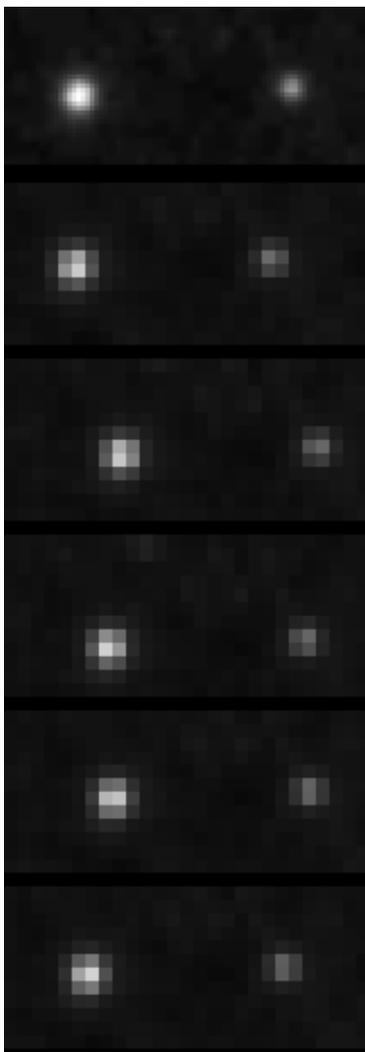
the Hubble Deep Field's WFPC2 camera.[1] The data collected using the wide-field camera on the Hubble are *undersampled*. The optics of the Hubble Space Telescope are able to resolve a lot more detail than the size of the CCD pixels will allow.

The image on the middle of this page is a demonstration of this. Here, I have used the [SkyView](http://skyview.gsfc.nasa.gov/cgi-bin/titlepage.pl) site (http://skyview.gsfc.nasa.gov/cgi-bin/titlepage.pl) to take a frame from the Digitized Sky Survey of M13 and cropped around two stars. The original image was taken at 1 arc-second per pixel and is shown in the top panel. Despite being magnified considerably in this image, the stars are both fairly round and smooth. They do not look very blocky or "pixelated" and can be said to be reasonably *well-sampled*. That is, we are using pixels small enough to accurately depict their shape without using many excess pixels.

The five images below this are all undersampled versions of the top image and simulate what these stars would look like if we used a CCD detector with pixels three times as wide and three times as high as the CCD detector used in the top image. Each of these images looks blocky and pixelated as a result.

Each of these five images looks slightly different, however, and it is this difference that is the key to *Drizzle*. In each of these five frames, the high-resolution image was shifted slightly (otherwise known as its location being "dithered") prior to being converted into the low-resolution image.[2] Thus, the fact that the stars move a bit from one image to another should come as little surprise.

The key insight behind *Drizzle* comes from a second difference we can see in the stars. Look closely and you'll see that each star *looks different* across the frames. The exact pattern of lighter and darker pixels used to depict each star varies from frame to frame as the star moves slightly. This variance

---

[1] A.S. Fruchter, & R.N. Hook (1998, 2002). Drizzle: A Method for the Linear Reconstruction of Undersampled Images. Available at: http://xxx.lanl.gov/pdf/astro-ph/9808087

[2] Note, the high-resolution image was not re-sampled in this process – technically, this was done by shifting the coordinate systems during the down-sampling into the lower-resolution images (performed in Matlab, by MathWorks).

Stacking Images with Drizzle *continued*

is not noise, but an accurate representation of the star as it is centered on different parts of the CCD array – sometimes dead-center over a pixel, sometimes slightly to the left of center of a pixel, sometimes slightly above and a little to the right of the center of a pixel, etc.

The fact that undersampled stars look blocky cannot be avoided. What Fruchter and Hook demonstrated in *Drizzle* was that we can exploit the fact that small motions of the stars – even fractions of pixels – make the stars look blocky in different ways each time to reconstruct higher resolution images. That is, we can use the "dither" or random noise in the *location* of the stars with respect to the CCD array in conjunction with many images to reconstruct an image that is higher resolution than our CCD array itself.

The Method

Here is a diagram from [Andy Fruchter's page](#) on Drizzle showing the technique. We first begin with a set or stack of original images. Each pixel in an original image is first reduced (from the red pixel size to the blue pixel size – this is the "pixfrac" or "pixel reduction factor" term). Like other stacking techniques, the original images are translated, rotated, and scaled as needed to align them onto an output image (the "Geometric transformation"). Here, this is done in a somewhat more complex way than typical stacking ever.

Coarse input pixel grid

Fine output grid

Geometric transformation
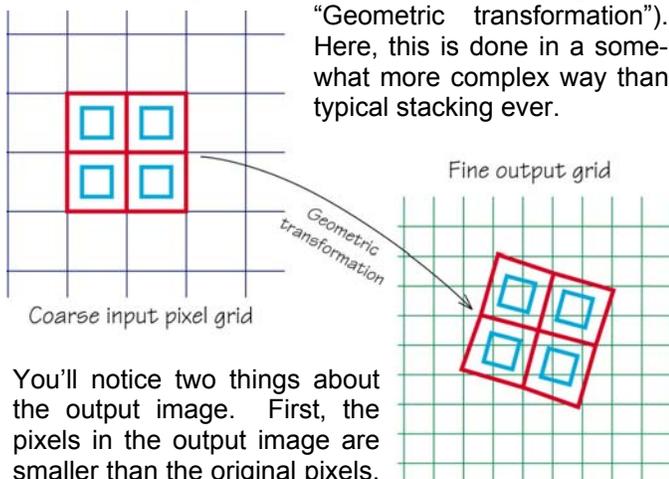
You'll notice two things about the output image. First, the pixels in the output image are smaller than the original pixels. This is the "up-sample factor" how – how much higher the output resolution is than the input resolution (aka how much bigger your output image is than your input image).

Second, these pixels will usually not line up perfectly with a pixel in the output image. Key to Drizzle is the

fact that these "drops" (the blue pixels) fall onto the output pixels according to the degree of overlap. The value dropped by a blue pixel (e.g., the intensity of a star picked up by a pixel in the original image) will fall a lot on one output pixel, a little on another, less still on a third, etc. Think of the output grid as having little wells (or locations on a very small ice-cube tray) for each output pixel and the water drop landing on a spot that hits multiple wells. The drop breaks up into smaller droplets that then fall into a number of wells as it hits the edges that separate the wells from each other. Therefore, as multiple images rain pixels down onto the output image, the output pixels fill up to the degree that input pixels line up with the output pixels.

Once all the drops from all your images have "Drizzled" onto the output image, it's a simple matter of computing the average intensity dropped onto that pixel, weighted by how big each droplet was. In the end, you'll have an image with a finer resolution than the raw frames you began with.

Results

That's the theory at least. But does it work? Given that the STSI uses this for Hubble shots is a pretty good indication that yes, in fact, the technique does work and, providing that you have enough data and that you don't ask it to make a silk purse out of a sow's ear, you can do quite well. Here, (image on following page) we have a cropped sample using that same shot of M13 as the reference image. The top panel shows the original image from the DSS. I took this image and created 50 undersampled versions of it by shifted it slightly each time and resampling the shifted images down to approximately half the original size.[3] The second panel shows what a

---

[3] Undersampling reduced images to 55% using bilinear interpolation and followed full pixel shifts of the reference image by 0-10 pixels in X and Y (performed in Matlab, by MathWorks). Stacking was performed in Nebulosity (Stark Labs) for both the traditional translation and Drizzle (pixel reduction factor = 0.4; up-sample factor = 2; atomizer = 3). The stacked images were resampled at 300%, 540%, and 270% using nearest-neighbor interpolation to equate image scale and accurately reflect the quality of each for display purposes (PhotoShop, by Adobe).

Stacking Images with Drizzle *continued*
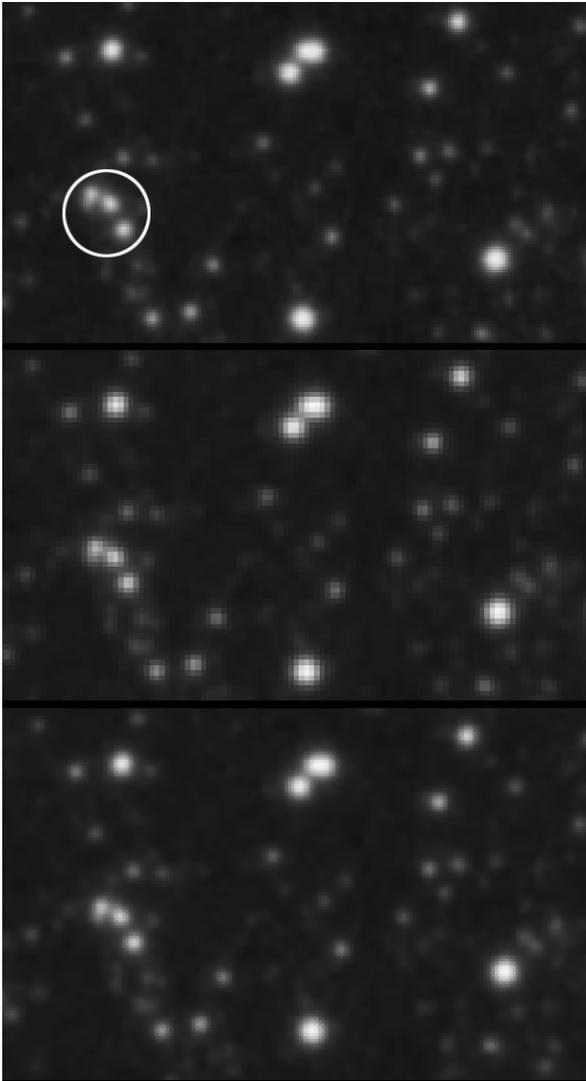
traditional stack looks like (scaled back up to match sizes on the screen). Note how the stars appear fairly blocky and how the region highlighted in the top panel shows no evidence of the fourth star. The bottom panel shows the same 50 frames aligned and stacked using Drizzle. While this is not a perfect reproduction of the original image, it clearly possesses a lot more detail than the traditional technique.

I mentioned earlier that Drizzle cannot turn the pro-

verbial sow's ear into a silk purse. If you've taken a half-dozen 640x480 webcam images with a 400 mm focal length telescope aimed at a small target like the Ring Nebula, don't expect Drizzle to turn this into a multi-megapixel image with sub-arcsecond resolution that pulls out fine structure in the Ring. It won't. It can make the image better, but it has its limits.

Try to do this and you'll quickly see "holes" left in the output image where no droplet fell or where not enough droplets fell to accurately estimate the true signal in that pixel. Recall that the output image's pixels are smaller than the input images pixels by the amount that you're up-sampling the image. If you up-sample the image too much or you don't have enough frames with enough motion to ensure that all output pixels have a reasonable amount of signal drizzled into them, you will end up with the kind of checkerboard pattern shown here.

Downsides

Drizzle works with black and white images, color images, images taken on an equatorial mount that only need shifting between frames, images taken on an alt-az mount that require shifts and rotations, and even images taken on different scopes that require shifts, translations and scales to adequately align.[4] Drizzle is a general-purpose stacking technique that recovers lost resolution during the stacking process. Why don't we all use Drizzle all the time?

1. Software support: Your favorite image processing software may not include the Drizzle algorithm. Envisage (packaged with Meade DSI cameras), Iris, Nebulosity, and Registax all include at least basic Drizzle support (images here were done in Nebulosity).

2. Time: Drizzle is not the fastest algorithm out there. In fact, it can be quite slow, especially if your camera already produces large files. Remember that if you "up-sample" your image by 2x, you need 4x the memory (it's 2x as big in each direction) and you're doing a lot more than simple translation and rotation. For each pixel in each image, you need to figure out exactly how much of that pixel, once moved and rotated overlaps each output pixel. This takes time.

3. Undersampling: Are your images actually undersampled? In many cases, especially if you're shooting wider-field shots, they are. But, if you're shooting your 8" SCT at 1 arcsecond per pixel to grab that planetary nebula, you're images aren't suffering from undersampled stars. Driz-

[4] Note, not all implementations of Drizzle support all of these features.

Stacking Images with Drizzle

zle can't help you pull out detail that's lost from tracking errors or seeing-induced blur.

4. Dithered locations: Drizzle assumes that the stars move in your image a bit. If you've got your autoguiding down so well that the stars don't even move half a pixel or so, Drizzle has no chance of working its magic. (If you do have such a setup, investigate intentional dithering of your guiding location to give Drizzle the motion it needs).

5. Number of frames: If you're the kind of imager who takes a small handful of very long exposures, Drizzle may not be your best choice. If you use it, you may need to use larger "drops" (higher values of "pixel fraction") and up-sample less aggressively. Drizzle only works if it can use many movements of the image to recover what was lost. If you don't give it enough information, you will get the checkerboard or "holes" pattern described above.

Conclusions

For many of us, the downsides are not very large as we typically take many individual frames, and often shoot at several arcseconds per pixel. While many amateurs guide their telescopes to reduce the mount's errors, unless you're guiding so well that you don't even bother to align your raw frames during stacking (which you could and should do if your tracking is actually perfect), there is motion enough between frames to drive Drizzle. Finally, if the computer needs to chunk away for a long time to process the images, we have a nice excuse to get up for a cup of coffee and stop staring at the computer screen for a bit.

Thus, Drizzle is an exciting technique to add to our arsenal. By cleverly exploiting one weakness in our images (motion between frames and the need to take many individual images), we can substantially address a second weakness (lost resolution due to pixelization) and sharpen our images. All without buying a new scope or camera! ♦

By day, Craig Stark, Ph.D. is a professor of Cognitive Neuroscience at Johns Hopkins University. By night, he is an amateur astrophotographer and developer of software for astrophotography Stark Labs
http://www.stark-labs.com