

How to co-add images? I. A new iterative method for image reconstruction of dithered observations

Lei Wang and Guo-Liang Li

Purple Mountain Observatory, Chinese Academy of Sciences, Nanjing 210008, China; leiwang@pmo.ac.cn,
guoliang@pmo.ac.cn

Received 2016 December 22; accepted 2017 June 1

Abstract By employing the previous Voronoi approach and replacing its nearest neighbor approximation with *Drizzle* in iterative signal extraction, we develop a fast iterative *Drizzle* algorithm, named *fiDrizzle*, to reconstruct the underlying band-limited image from undersampled dithered frames. Compared with the existing *iDrizzle*, the new algorithm improves rate of convergence and accelerates the computational speed. Moreover, under the same conditions (e.g. the same number of dithers and iterations), *fiDrizzle* can make a better quality reconstruction than *iDrizzle*, due to the newly discovered High Sampling caused Decelerating Convergence (HSDC) effect in the iterative signal extraction process. *fiDrizzle* demonstrates its powerful ability to perform image deconvolution from undersampled dithers.

Key words: techniques: image processing — methods: observational — stars: imaging — planets and satellites: detection — gravitational lensing

1 INTRODUCTION

All imaging processes involve a limitation related to resolution of the equipment. In practice, the number of detectors is limited, thus the sampling is limited. Since spatial frequencies in an astronomical image are strongly limited by the optics of the telescope, the band is limited. For economical or other considerations, e.g. to cover a wide field in each exposure, the detector sampling sometimes cannot reach the Nyquist (or critical) sampling of the optics in a telescope. Therefore the detector often collects a set of undersampled data. In fact, an undersampled detector inevitably blurs the details in its sampling interval. This blurring effect is so-called *aliasing*. When the sampling process is executed by a CCD (or CMOS) pixel matrix (via a digitizer), the effect of *aliasing* is expressed as *pixelation*.

In order to restore details lost in *pixelation* or *aliasing*, researchers have proposed increasing the sampling rate by increasing the number of exposures of the same field (but with different shifts, i.e. dithered frames). Thus the remaining question is how to reconstruct the signal from dithered frames. Many methods have been developed such as *interlace* and *shift-and-add*, which can ob-

tain a high resolution result and reduce the pixelation to some extent, but these results are still far from excellent *anti-aliasing*. By taking advantage of *interlace* and *shift-and-add*, Fruchter & Hook (2002) improve on the previous *shift-and-add*, named *Drizzle*. However, like the method mentioned before, *Drizzle* does not enhance the anti-aliasing function though it has a better performance than previous works in reducing noise and increasing accuracy. In fact, *Drizzle* generates a flux averaged image on a high resolution grid, thus producing a blurred, contrast reduced appearance. Based on a non-parametric method called *kernel regression*, which takes both the relative spatial and radiometric distances of nearby pixels into account, Takeda et al. (2006) developed an improved method named *super-Drizzle*, which can reconstruct a high quality image compared with *Drizzle*. *super-Drizzle* has good performance in de-convolving the pixelation to some extent. However, *super-Drizzle* is more likely to be applied for image denoising and interpolation. It sensitively depends on the number of dithered frames and parameter selection. Thus it is difficult for *super-Drizzle* to obtain a higher contrast image than *Drizzle* when the dithers are not enough. By replacing the value of nearest neighbor with that of *Drizzle* in an iterative

Voronoi approximation [initially developed by Werther (1999), Gröchenig & Strohmer (2001)] and introducing the *oversampling - low pass filtering - interpolating* process to the image co-adding procedure, Fruchter (2011) upgraded the previous Voronoi approximation to yield *iDrizzle*. *iDrizzle* was developed largely for creating accurate images of objects with unresolved or nearly-unresolved components. With the help of iterative signal extraction and low pass filtering in the frequency domain, *iDrizzle* deconvolves the pixelation of undersampled features (with high signal-to-noise ratio, SNR) much better than *super-Drizzle* on small scales. However, *iDrizzle* has an oversampling, filtering and interpolation process, which dramatically increases the amount of computation.

In this paper, compared with the previous *iDrizzle* method, we improve on the effectiveness and computational speed, by introducing a program called fast *iDrizzle* or *fiDrizzle*. *fiDrizzle* can accelerate the computation and improve image quality. After describing the *fiDrizzle* algorithm and analyzing the theoretical aspects of its mechanism in Section 2, in Section 3 we illustrate the results to make a visual (Sect. 3.1) and a quantitative comparison (Sect. 3.2). We present the computational complexity analysis in Section 4 to show how the new algorithm can accelerate the computation. The algorithm's dependency on the number of dithered frames and iterations is shown in Section 5. Finally, the discussion and conclusion are provided in the last section (Sect. 6).

2 THE FIDRIZZLE ALGORITHM

Drizzle is superior in computational speed. However, its intrinsic filter (pixelation) removes high frequency information on small scales. *iDrizzle* can reconstruct details on small scales to some extent at the cost of a very large amount of computation and a huge volume for the output file. In order to maintain the advantages of both *Drizzle* and *iDrizzle* and reduce their weaknesses, we develop *fiDrizzle* to improve image co-adding technology. The algorithm is described in the following steps:

Step 1: Apply *Drizzle* to the K dithered images of a field, $\{I^1, I^2, I^3 \dots I^K\}$, onto a high resolution grid one needs¹, e.g. critical sampling grid in the below examples of this work, to produce the image D_1 , which is the first approximation to the true image. The 1 in the subscript of D indicates that this is the first iteration.

Step 2: Map the first approximation image, D_1 , back to the frames of the original individual images by a series

of mimic observations², which produces a set of approximations to the original images, A_1^m .

Step 3: Subtract the mimic observations from the corresponding original images to produce a series of residual images $I_1^m = I^m - A_1^m$.

Step 4: Return to the first step and now apply *Drizzle* to the set of residual images $\{I_1^1, I_1^2, I_1^3 \dots I_1^K\}$ to produce the image D_2 which is the difference between the true image and the previous approximation.

Step 5: Continue as before with one modification, at **Step 3**, in the N th iteration, $A_N = \sum_{j=1}^N A_j$, until the residuals are dominated by noise.

After the iterations are complete, one can regard the final approximation as the best fitting image to the true one. Comparing with the previous *oversampling - low pass filtering - interpolating* process in the steps of *iDrizzle*, we basically (i) remove the low pass filtering process; (ii) directly co-add the observation frames to an (at most) critically sampled grid instead of an oversampled one; (iii) therefore, the final *sin c* interpolation is not necessary. The improvement accelerates the computation significantly. Compared with the Voronoi approximation, we just replace the nearest neighbor approximation with the *Drizzle* result at each iteration.

In order to figure out the difference between *iDrizzle* and *fiDrizzle* in theory, we will compare their result under the same resolution grid, e.g. critical sampling. It means that *fiDrizzle* directly samples the original image to critical output, while *iDrizzle* undergoes an *oversampling - low pass filtering - interpolating* process to obtain the same resolution. Let T be a signal in two dimensional space. After being undersampled by equipment, e.g. a telescope, one gets an image I with a dither shift ds (relative to the output target grid, including position and angle shift),

$$I = T \otimes G, \quad (1)$$

where G represents all the combined effects from the true signal T to the equipment, such as seeing, PSF, pixelation, CCD distortion, etc. Symbol \otimes represents the convolution operator. Now we resample I to an oversampling D_1^{OS} and a critical sampling D_1^{CS} , and in the meantime take the dither shift ds into account.

$$\begin{aligned} D_1^{\text{OS}} &= I \otimes P^{\text{OS}}, \\ D_1^{\text{CS}} &= I \otimes P^{\text{CS}}, \end{aligned} \quad (2)$$

where P^{OS} or P^{CS} is the resampling matrix which includes a pixelation effect. Also, P^{CS} can be obtained by $P^{\text{CS}} = P^{\text{OS}} \otimes \sin c(P^{\text{OS}} \rightarrow P^{\text{CS}})$, where

¹ In general, the sampling should not go beyond the Nyquist sampling (critical sampling), which is different from the indispensable oversampled grid in *iDrizzle*.

² Mimicking the dither; all exposure conditions are identical to those of the original dithered observations.

$\sin c(P^{\text{OS}} \rightarrow P^{\text{CS}})$ is lossless interpolation of the signal from oversampling to critical. Therefore D_1^{CS} is equal to $D_1^{\text{OS}} \otimes \sin c(P^{\text{OS}} \rightarrow P^{\text{CS}})$. Here D_1^{OS} or D_1^{CS} is the very result of *Drizzle*. Following the iterative reconstruction steps, we mimic the real observation by down sampling the target grid D_1^{OS} and D_1^{CS} to the I grid, thus getting the first approximations A_1^{OS} and A_1^{CS} to the original image I respectively. Then the difference (or residual image) from the original observation I can be expressed as:

$$\begin{aligned} I_1^{\text{OS}} &= I - A_1^{\text{OS}}, \\ I_1^{\text{CS}} &= I - A_1^{\text{CS}}. \end{aligned} \quad (3)$$

Since A_1^{OS} has higher sampling than A_1^{CS} , correspondingly the residual image I_1^{OS} possesses less power than I_1^{CS} , especially at the low frequency end. In an extreme case, if one oversamples the original image I to an infinitely high resolution grid, D_1^{OS} will keep an intact I , leading to $A_1^{\text{OS}} = I$, and thus $I_1^{\text{OS}} = \mathbf{0}$. After the first iteration, we have the second approximations to the true image, D_2^{OS} and D_2^{CS} :

$$D_2^{\text{OS}} = D_1^{\text{OS}} + I_1^{\text{OS}} \otimes P^{\text{OS}}, \quad (4)$$

$$D_2^{\text{CS}} = D_1^{\text{CS}} + I_1^{\text{CS}} \otimes P^{\text{CS}}. \quad (5)$$

Following the last step of *iDrizzle*, one may ($\sin c$) interpolate the oversampled approximation to the critical sample grid with

$$\begin{aligned} D_2^{\text{OS}} &\otimes \sin c(P^{\text{OS}} \rightarrow P^{\text{CS}}) \\ &= D_1^{\text{OS}} \otimes \sin c(P^{\text{OS}} \rightarrow P^{\text{CS}}) \\ &\quad + I_1^{\text{OS}} \otimes P^{\text{OS}} \otimes \sin c(P^{\text{OS}} \rightarrow P^{\text{CS}}) \\ &= D_1^{\text{CS}} + I_1^{\text{OS}} \otimes P^{\text{CS}}. \end{aligned} \quad (6)$$

Comparing Equation (6) to Equation (5), we find that for the same coarser (here critical) sampling case, iterative reconstruction on critical sampling gains more power (and hence more details) than that on oversampling, which means for the same number of iterations the rate of convergence for iterations is decelerated by the high sampling (here oversampling) case. We call this effect the *High Sampling caused Decelerating Convergence* (HSDC)³ in the iterative signal extraction process. A simple simulation that serves as proof is provided in Appendix A of this paper. In fact, oversampling leaves more (original) pixelation effect in the final result than critical sampling. Therefore, *fiDrizzle* works more effectively in pixelation deconvolution than *iDrizzle*, even if

³ It is easy to prove the HSDC effect in other dimensions for any number (≥ 1) of iterations and dithered observations.

we ignore the inserted filtering in the frequency domain in *iDrizzle*. In the next section, we provide more examples to test the HSDC effect and check the validity of the *fiDrizzle* algorithm.

3 THE RESULTS

We show visual and quantitative comparisons of all the three above mentioned image co-adding methods, *Drizzle*, *iDrizzle* and *fiDrizzle*, in the following subsections. According to sampling theory, if one wants to double the spatial resolution of the digital signal which is originally extracted from an analog signal, he (or she) should double the sampling frequency directly or double the number of observations (in different positions) but keep the sampling frequency unchanged, i.e. double the number of dithers (which is also the sense of dither). Therefore, in order to totally restore the signal from a set of $\frac{1}{2} \times \frac{1}{2}$ undersampled observations, one should have at least four dithered frames to construct a critically-determined (or over-determined, which needs more dithers) system. Since this work mainly compares the co-adding methods on the critical sampling, in the following analysis, the systems are all over-determined for the critical sampling, but under-determined for the over-sampling, e.g. *iDrizzle*. We have checked a similar result when the system is over-determined for both critical and over sampling. Theoretically, an over-determined system can significantly reduce degeneracy from the effect of the random dithers (not well-placed dithered frames). In order to mimic a series of dithered frames, the true image is dithered to several undersampled frames by introducing random shifts, rotations and CCD geometric distortions of 0.1%.

3.1 A Visual Comparison

We use the well-known picture of Lena (with size 512×512 pixel, assuming this resolution meets the Nyquist (critical) sampling, i.e. the sampling frequency is at least twice the highest spatial frequency, which can be resolved by a telescope) to check the performance of the three co-adding methods. In Figure 1, the true image (left panel) is binned to five dithered frames with a $\frac{1}{2} \times \frac{1}{2}$ lower resolution (right panel, undersampled) than critical. Considering rotation and position shift, the dithered frames have a pixel size of at most $512/\sqrt{2} \times 512/\sqrt{2}$ (the critical sampling pixel size, i.e. $256/\sqrt{2} \times 256/\sqrt{2} \simeq 181 \times 181$ original pixels), in which every pixel in the frame will not go beyond the region of the true image. In the dithered frame one can identify the effect of aliasing in regions with rich details, pixelation at



Fig. 1 An image of Lena at different resolutions. One of the dithered images (*right panel*, which is usually regarded as the original input in image co-adding) has $\frac{1}{2} \times \frac{1}{2}$ lower resolution than the true image (*left panel*). The effect of aliasing significantly smooths details on small scales. Details in 2×2 true image pixels (corresponding to one original pixel) are averaged out to a single value.



Fig. 2 Reconstructions from three different methods. The upper left image is the true one. The upper right panel is from *Drizzle*, the lower left from *iDrizzle* and the lower right from *fiDrizzle*. Both *iDrizzle* and *fiDrizzle* are applied in five iterations. The three reconstructions have a 2×2 higher resolution than the original and thus have the same resolution as the true one.




Fig. 3 Residuals between the reconstructions and the true case (*upper left panel*). Upper right is for *Drizzle*, lower left for *iDrizzle* and lower right for *fiDrizzle*.

the edge of the hat, blurred stripes in the body of the hat, and eyes and lips dimming. Such a mosaic-like image also exhibits loss of contrast and gray levels. In short, the effect of aliasing makes the image blurred and pixelized.

Figure 2 shows three reconstructions from the identical set of the (five) dithered frames: the upper left image is the true one, the upper right one is reconstructed by *Drizzle*, the lower left one is from the $\sin c$ interpolation of the oversampled⁴ *iDrizzle* result and the lower right one is produced by this work, i.e. *fiDrizzle*. Following the strategy described in Appendix B, we choose a proper mask function with $r_f = 230$ for the filtering steps in *iDrizzle*⁵. In the lower panels, both *iDrizzle* and *fiDrizzle* are applied in five iterations. Obviously, the *Drizzle* result (the upper right) is better than the dithered image as shown in the right panel of Figure 1. However, compared with the result from *iDrizzle*, *Drizzle* is not able to restore details at the level of a few pixels. There is an inherent filter (an averaging effect introduced by pixelation) convolved by the *Drizzle* mechanism, which results in high frequency information missing. That is the reason why the upper right image looks smoothed and blurred. Due to repeated signal extraction from the residual image, both *iDrizzle* and *fiDrizzle* co-add images much better than *Drizzle* does. Although the two lower images have a similar appearance, after a careful eyeball check, the right one looks sharper with higher contrast than the left. In addition, *fiDrizzle* generates more stripes⁶ in the body of the hat, as well as sharper contrast in the hair and eyelashes, and thus yields better image quality than *iDrizzle*.

Obvious differences are shown in the residual images (Fig. 3). In Figure 3, the upper left one is the same as in Figure 2, but the rest are differences between the reconstructions (generated by the three algorithms) and the true one. Three residual images are scaled to the same scope, then have the same color bar scale. For the portrait of Lena in the shadowed area, the more recognizable the figure is, the more signal it loses in image reconstruction. Evidently, the *Drizzle* algorithm loses lots of information, then shows a significant portrait in the residual image. *iDrizzle* misses a few details in high frequency. As a result, it leaves some features in regions with rich detail, e.g. hair, eyelashes and stripes in the hat. However, the residual from *fiDrizzle* is almost unrecognizable, which turns out to be the best fitting to the true case among the

three results. If zooming in the lower panels and focusing on the sharp transition edges, one can find *iDrizzle* introduces a ringing artifact, which appears as *ghosts* near transients.

Furthermore, we try to figure out how the visual difference affects their power in the frequency domain. Here we introduce a reduced power spectrum (RPS, reduced to one dimension) to analyze the power left in the above three residual images. In order to avoid the regions in which pixels are not totally covered by all dithered frames, we select an all-covered area which is generally 1/4 times the true case, i.e. a pixel size of 256×256 (critical sampling pixel) and has the same center as the critical sample. We define the RPS as the radial power distribution in the fast Fourier transform (FFT) image of the all-covered area. Therefore, in the frequency direction the RPS has 128 pixels. In Figure 4, the black line stands for RPS of the true case. Other colors are the RPS (lower is better) of the three residuals between the true case and three reconstructions: green for *Drizzle*, blue for *iDrizzle* and red for *fiDrizzle*. Here one can find that most high frequency power is left in the three residuals. As expected, *Drizzle* is the worst one. Due to the oversampling mechanism and low-pass filtering, *iDrizzle* performs a little better than *fiDrizzle* at the high frequency end. However, *iDrizzle* loses much more power in low and medium frequencies than *fiDrizzle* because of the HSDC. This is the reason why Lena reconstructed by *fiDrizzle* has a high contrast level in Figure 2 and why the residual from *iDrizzle* in Figure 3 is still recognizable.

3.2 The Quantitative Comparison

In weak gravitational lensing astronomy, people usually calculate the lensing effect by measuring the shapes of background galaxies and comparing them with a randomly oriented sample of galaxies. Its final result sensitively depends on the accuracy of shape measurement, which means a high fidelity image co-adding method can significantly improve the SNR of a weak lensing signal, thus enhancing the accuracy and reliability of the result. We extract a spiral galaxy image with a resolution of 512×512 from the *Hubble Space Telescope* (HST) observation HST_jc1g03010_drc.fits at R.A. = 195.01238 deg and Dec = 28.023106 deg as the true picture to test. Then the true image is dithered to 10 frames (down sampled to $\frac{1}{3} \times \frac{1}{3}$ lower resolution) by applying random shifts, rotations and CCD geometric distortions to mimic the undersampled dithered observations. Note that here *iDrizzle* has an oversampled grid with resolution 4×4 times the critical case. During the re-

⁴ The sampling of the oversampled image is 2×2 times the Nyquist sampling.

⁵ Note that there are a total of 256 pixels from the FFT image center to its edge. Therefore, $r_f = 230$ is close to the Nyquist frequency.

⁶ Strictly, this is due to the improvement in contrast level.




Fig. 4 The RPS for residuals between reconstructions and the true image. The RPS for *Drizzle* is in green, *iDrizzle* in blue and *fiDrizzle* in red, while the black line is the RPS of the true image (NOT the residual).

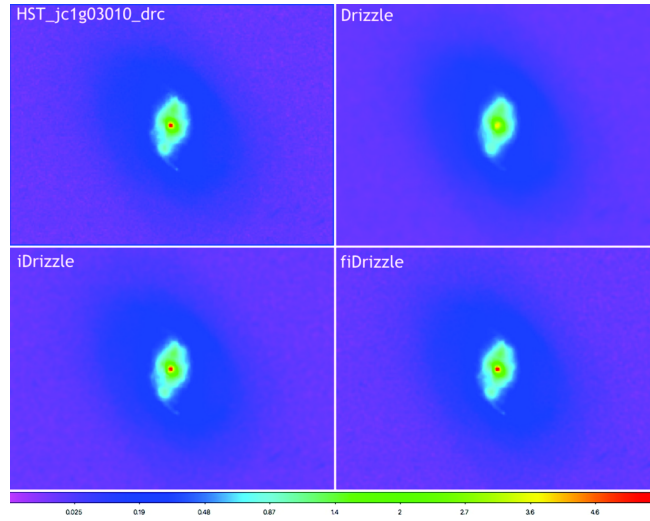


Fig. 5 Another image reconstruction test, using *HST* data (with image center at R.A. = 195.01238 deg and Dec = 28.023106 deg). The layout of panels is similar to Fig. 2.

construction, six iterations are executed in both *iDrizzle* and *fiDrizzle*. We plot the three kinds of reconstructions in Figure 5 with the same layout as in Figure 2. The color stands for the flux received by the pixels that were part of the observing equipment and is already scaled to the same scope. Also we show the residual plot in Figure 6 with the same representation and layout as in Figure 3. Similar to the result in Section 3.1, both *iDrizzle* and *fiDrizzle* recover better image quality than *Drizzle*. However, the visual difference between *iDrizzle* and *fiDrizzle* is not significant. So, we investigate the flux at the pixels that satisfy $X = Y$ in the four panels of Figure 5.

In Figure 8, the flux profile, which is normalized to the flux of the central pixel of the true image, is plotted in the upper panel: the true one is in black, *Drizzle* in green, *iDrizzle* in blue and *fiDrizzle* in red. The lower panel shows the flux profile of the $X = Y$ pixels in the residuals displayed in Figure 6. Here the size of 3×3 critically sampled pixels equals that of one original pixel. Evidently, *fiDrizzle* provides the best fitting result to the true case, especially at the center. Figure 7 shows a similar result as that in Figure 4: *fiDrizzle* is the best one in low and medium frequencies, but also leaves a few high frequency noises in the reconstruction. However, for shear measurement of weak gravitational lensing, low frequency information plays an important

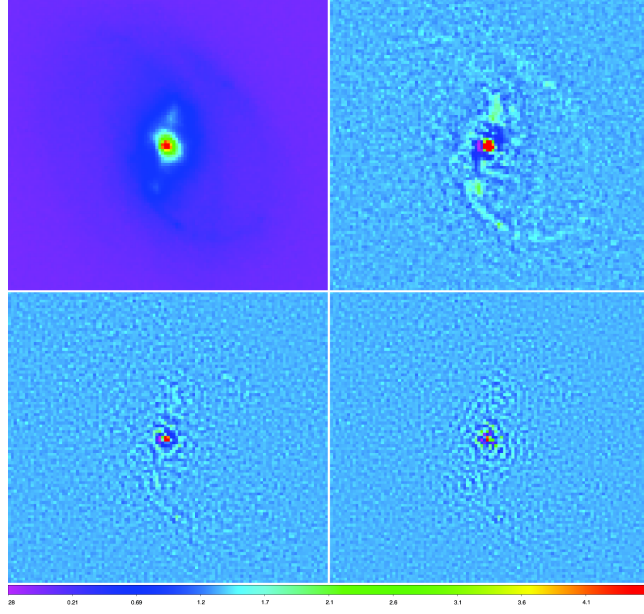


Fig. 6 Residual maps for the *HST* image reconstructions. The layout of panels is the same as in Fig. 3.




Fig. 7 The residual RPS for the *HST* image reconstructions. The color representation of lines is the same as that in Fig. 4.

role. Following Hirata & Seljak (2003), the ellipticity of an object is defined as

$$\begin{aligned} e_+ &= (M_{xx} - M_{yy}) / (M_{xx} + M_{yy}), \\ e_\times &= 2M_{xy} / (M_{xx} + M_{yy}), \end{aligned} \quad (7)$$

where M_{ij} represents the moments (see Hirata & Seljak (2003) for details). The spin-2 tensor $\mathbf{e} = (e_+, e_\times)$ is the so-called ellipticity tensor. In order to avoid the problem of divergence, a circular Gaussian weighting function with a weight radius of r_w is convolved into the four images in Figure 5 before the measurement. We then plot the ellipticity tensor \mathbf{e} of the source as a function of

weight radius r_w and show them in Figures 9 and 10. The color definition of line types is the same as that in Figures 8 and 4. Since only 39.3% of the weight is located in an area with radius r_w , but 86.5% is in $2 \times r_w$, here we use $2 \times r_w$ as the variable, which means if one faces a uniformly illuminated source, the flux from pixels in $\leq 2r_w$ contributes 86.5% of the total to the measurement.

From low to high frequency, *Drizzle* has no advantage compared with *iDrizzle* or *fiDrizzle*. This reflects how important pixelation deconvolution is in the image co-adding process. In the ellipticity measurement, power




Fig. 8 The flux and residual profile of the three reconstructions at the $X = Y$ pixels. Lines have the same color definition as in Fig. 4. Note that the length of three pixels here corresponds to one original pixel.




Fig. 9 Measuring the ellipticity component e_+ as a function of weight radius r_w . The error bars are estimated from the noise in the image. Line types are the same as those in Fig. 8.

at low frequency determines the signal, while noise at high frequency mainly affects the scatter of the final result. In Figure 9, *fiDrizzle* behaves better than *iDrizzle* from small scales to large, and has the lowest systematic error among the three reconstructions at large scales. Note that gravitational lensing is very sensitive to systematic error. It turns out that the HSDC effect is the main reason that prevents *iDrizzle* from obtaining enough low frequency power and reducing the systematic bias in reconstruction. However, there is little difference between *iDrizzle* and *fiDrizzle* in the e_x component of ellipticity in Figure 10. In this case, we find the disadvantage of

fiDrizzle at high frequency does not badly affect the shear measurement because most of the high frequency noise is averaged out so as to be negligible at large scales.

4 THE COMPUTATIONAL COMPLEXITY

There is no doubt that *Drizzle* runs much faster than the other two algorithms because it does not need more iterations. The computation of each mimic observation [or `blot` program in Fruchter & Hook (2002)] is similar to that of *Drizzle*. Then *fiDrizzle* costs $2N$ times the computational complexity of *Drizzle*, which depends on the




Fig. 10 Similar to the profile in Fig. 9 but for another ellipticity component e_x .




Fig. 11 Residuals from *fiDrizzle* for different numbers of dithers. The upper left panel is for one dither, the upper right panel for two dithers, four frames are in lower left panel and five frames are in the lower right.

number of iterations N . In Figure 2, the true image has a size of 512×512 pixel, thus each original frame (observation) has a pixel size⁷ of $256/\sqrt{2} \times 256/\sqrt{2} \simeq 181 \times 181$. Note that the workload of *Drizzle* depends not only on the number of total original pixels, $5 \times 181 \times 181$, but also on the resolution of the output grid (for this ex-

⁷ This size ensures that rotated and shifted dithered images (or field of view) cannot severely exceed the region of the true image.

ample, the critical sampling is 2×2 times higher than the original grid). So, *Drizzle* costs at least $T(n) = 5 \times 181 \times 181 \times 2900 \simeq 470\,000\,000$ computations⁸, and thus has a computational complexity $O(2000n^2)$ (set $n = 512$). After five iterations, *fiDrizzle* has a total com-

⁸ Here we convert the surface integral into a line integral via the Green formula. After a series of polygon clipping operations and line integral computations, it costs ~ 2900 floating-point calculations in each original pixel.

putational complexity $O(20\,000n^2)$, due to both *Drizzle* and the mimic observation process in each iteration. As for *iDrizzle*, there are three expenditures in amount of calculation:

- First, *iDrizzle* requires an oversampled output grid which has a 2×2 (in this example) times higher resolution than the critical sampling. Then it results in a 4 times higher number of computations than *fiDrizzle*, i.e. a complexity $O(80\,000n^2)$.
- Second, the workload of oversampled image FFT, smoothing and inverse FFT are, at least, $T(n) = 5 \times 2 \times 1024^2 \times 2 \times \log_2(1024) \simeq 210\,000\,000$, i.e. the total computational complexity is $O(80n^2 \log_2 n)$.
- Third, the final sinc interpolation contributes an $O(64n^2)$ computation complexity⁹.

Compared with *fiDrizzle*, *iDrizzle* is mainly delayed by the oversampling strategy. The CPU time consumed by the filtering process is about a few percent of the total when the size of an oversampled image is about 1024×1024 . Combining the results above, we find that *iDrizzle* is not only decelerated in the rate of convergence (the HSDC effect) but also delayed in computational speed for the same reason — the oversampling strategy.

5 DEPENDENCY ON THE NUMBER OF DITHERED FRAMES AND ITERATIONS

In this section, we discuss how *fiDrizzle* or *iDrizzle* depends on the number of dithered frames K and iterations N . We run *fiDrizzle* to reconstruct the picture of Lena for one, two, four and five frames. Each reconstruction is performed in five iterations. $N = 5$ is a tradeoff between signal extraction and artifact reduction, such as the ringing artifact near sharp transitions, which is introduced by low pass filtering but is enhanced by successive iterations. The residual (the difference between the *fiDrizzle* reconstructions and the true image) plot is shown in Figure 11; the upper left panel is for one dither, the upper right panel is for two dithers, four frames are in the lower left panel and five frames are in the lower right. Obviously, the residuals are significantly reduced when we increase the number of co-added frames. Moreover, a strong argument can also be found in the RPS plot of the residuals in Figure 12. In order to compare *fiDrizzle* (in solid lines) with previous work, we also plot the results of *iDrizzle* (in dotted lines, with the same filter as Figure 2) in the RPS figure. Note that the color representation is totally different from the above figures: reconstructions for one, two, four and five frames are in purple, red, blue and

black respectively. Figure 12 shows that the quality of reconstruction strongly depends on the number of dithers when K is small. While the degree of this dependence decreases as K increases, one can also find that compared with *fiDrizzle* the advantage of *iDrizzle* on the high frequency end is diminished as the number of dithers K increases.

Now we fix the number of dithers K to be 5 and change the number of iterations $N = 0, 1, 3, 5$ for both *fiDrizzle* and *iDrizzle*. Here we only show the *fiDrizzle* reconstructed picture of Lena in Figure 13; the upper left panel is for zero iterations, namely the *Drizzle* result, the upper right panel is for one iteration, three iterations are executed in the lower left panel and five iterations are in the lower right. The efficiency of signal extraction is very high at the beginning several iterations, which results in the portrait of Lena vanishing quickly. Sufficient evidence is shown in Figure 14. Reconstructions for zero, one, three and five iterations are in purple, red, blue and black respectively (solid lines). As before, we also illustrate the results from *iDrizzle* with corresponding iterations and colors, but in dotted lines. Note that for the case of zero iterations with *iDrizzle*, we perform the filtering process after the first *Drizzle* step is complete, with no more signal extraction steps. From Figure 14, due to the HSDC effect we find that *fiDrizzle* converges more effectively than *iDrizzle* in low and medium frequencies. The low frequency difference between solid and dotted (same color) lines becomes large as the number of iterations N increases.

6 DISCUSSION AND CONCLUSIONS

The *oversampling - low pass filtering - interpolating* process is a standard industry practice for improving the SNR in analog to digital (A/D) signal transition and extraction. Naturally, this process is adopted by the previous work related to *iDrizzle*. Of course, there is no problem if one initially wants an oversampled reconstruction or the process does not involve iterative signal extraction from the residuals. However, the *oversampling - low pass filtering - interpolating* process and iterative signal extraction coexist in *iDrizzle*, which inevitably encounters the HSDC effect. As a result, compared with *fiDrizzle*'s direct sampling to the critical case, *iDrizzle* not only costs more computational resources but also converges ineffectively which leads to an inadequate reconstruction for low frequency signals and, eventually, affects the systematic errors in weak lensing shear measurement as described in Section 3.2. Briefly, in this work we reach some goals:

⁹ We use 16 points at each interpolation.




Fig. 12 The residual RPS of *iDrizzle* and *fiDrizzle* for different numbers of dithers. The reconstructions for one, two, four and five frames are in purple, red, blue and black respectively. *fiDrizzle* is shown in solid lines, while *iDrizzle* is in dotted lines.




Fig. 13 Residuals from *fiDrizzle* for different numbers of iterations. The upper left panel is for zero iterations, the upper right panel is for one iteration, three iterations are executed in the lower left panel and five iterations are in the lower right.

- We discover the HSDC effect in the iterative signal extraction process and mathematically prove its existence.
- For the same number of iterations, *fiDrizzle* converges more effectively than *iDrizzle*, especially at low and medium frequencies, thus obtaining a better quality reconstruction.
- Instead of oversampling the frames to a high resolution grid (*iDrizzle*), *fiDrizzle* directly samples the dithers to the critical resolution and omits the filtering and interpolation procedures, which finally saves more computational resources.

As mentioned before, *iDrizzle* can generate accurate images of objects with unresolved or nearly-unresolved components. *fiDrizzle* inherits this function as well if one co-adds the dithers in an oversampled grid at the very beginning. However, any features less than the scale of the maximum angular resolution of the optics are unbelievable, which are smoothed by the filter in *iDrizzle*, but retained in *fiDrizzle*. Nevertheless, it does not affect the photometry in both *iDrizzle* and *fiDrizzle*. In that sense, we do not include a comparison of unresolved features in the paper.




Fig. 14 The residual RPS of *iDrizzle* (dotted lines) and *fiDrizzle* (solid lines) for different numbers of iterations. Reconstructions for zero, one, three and five iterations are in purple, red, blue and black respectively.



Fig. 15 Reconstructions of one dithered frame (upper right) in five iterations for oversampling (lower left) and critical sampling (lower right) cases. The upper left is the true image.



Fig. 16 Residuals between the reconstructions (in Fig. 15) and the true image for oversampling (left) and critical sampling (right) cases.




Fig. 17 The residual RPS profile for different iterations in oversampling and critical sampling cases. One iteration oversampling is in green and critical sampling is in red. The result of five iterations for oversampling is in blue and critical sampling is in black.




Fig. 18 The residual RPS profile for *iDrizzle* with different filtering parameters: $r_f = 130$ is in red, $r_f = 150$ in purple and $r_f = 200$ in blue. The black line is for *fiDrizzle*.

It is worth mentioning that compared with *iDrizzle*, *fiDrizzle* displays its lack of reconstruction accuracy at the high frequency end. In our upcoming work, on one hand we improve *fiDrizzle* to enable it to restore a part of details at the high frequency, and on the other hand, we develop a totally new image co-adding method called Tessellated Simple Surface Fitting (TSSF), which can effectively balance pixelation deconvolution and noise reduction.

In the future, many new telescopes will start astronomical observation, e.g. NASA's Wide Field Infrared Survey Telescope (WFIRST), the European Space Agency (ESA)'s Euclid, the National Science Foundation (NSF) funded Large Synoptic Survey Telescope (LSST) and the Chinese Space Station optical Telescope (CSST). Huge amounts of imaging data will be generated by those telescopes. How to effectively and efficiently process these data will be an urgent re-

quirement. We believe that the *fiDrizzle* algorithm has some advantages and can make some contributions in astronomical image processing as long as undersampled dithers exist.

Acknowledgements We thank the anonymous referee for very helpful comments that greatly improved the presentation of this paper. This work is supported by the National Basic Research Program of China (973 program, Nos. 2015CB857000 and 2013CB834900), the Foundation for Distinguished Young Scholars of Jiangsu Province (No. BK20140050), the ‘Strategic Priority Research Program the Emergence of Cosmological Structure’ of the CAS (No. XDB09010000) and the National Natural Science Foundation of China (Nos. 11333008, 11233005, 11273061 and 11673065).

Appendix A: A SIMPLE EXAMPLE FOR TESTING THE HSDC EFFECT

Here we use one of the original images in Figure 1. The result is shown in Figure 15. The upper-left panel is the true image. After a dither observation one gets a $\frac{1}{2} \times \frac{1}{2}$ lower resolution image as the upper-right one (for convenience of comparison, here zoomed in 2×2). The lower-left panel is a reconstruction on a 2×2 higher resolution grid (oversampling) than the critical sampling and finally interpolated back to the critical sampling after a 5-iteration process, namely $P^{CS} \otimes \sum_{j=1}^5 D_j^{OS}$. The lower-right panel shows the critical sampling reconstruction, i.e. $\sum_{j=1}^5 D_j^{CS}$. Obviously, the original dithered frame (upper-right) has the worst quality among the four panels. Aliasing and blurring dominate the edge of high contrast and the stripes in the hat. If observing carefully one can find the lower-right image has higher gray resolution than the lower-left one (due to the HSDC effect). However, the lower-left image has a few advantages over the lower-right one in the high frequency end (see also Figs. 16 and 17).

The difference between the over (or critical) sampling reconstruction and the true image is shown in Figure 16, left for the oversampling residual, right for the critical sampling. The difference is very significant in visual appearance.

A similar result can be found in the frequency domain, shown in Figure 17. The definition of Figure 17 has been made in Section 3.1. Significantly, critical sampling is much more advantageous at the low frequency end, however, there is a little worse performance at the high frequency end than the oversampling mode. This example also conveys such a message that even if somebody has only one dithered image, he (or she) can get more details through the iterative reconstruction process only if the frame’s position information can be taken into account.

Appendix B: CHOOSE A PROPER FILTER FOR IDRIZZLE

A gentle taper in Fourier space can suppress the ringing effect. However, it also suppresses the low frequency which is necessary in image reconstruction. Following Fruchter (2011), a similar filter, which has a sharp cut-off near the Nyquist frequency, is adopted in this paper to balance the ringing suppression and the low passband requirement. Here a circular mask (in Fourier space) that falls from 90% to 10% transmission over a width of 0.1 times Nyquist frequency is used. Figure 18 is an example for how to choose a proper filter for the *HST* image reconstruction case in Section 3.2. Here we keep the transmission intact, but tune the characteristic radius of the circular filter r_f (a distance from the center of the FFT image to the inflexion of the transmission). Finally, a relatively good case with $r_f = 150$ is selected. In order to ensure its validity, we also check it by visual and quantitative comparison with other cases, e.g. $r_f = 130$ or 200.

References

- Fruchter, A. S., & Hook, R. N. 2002, *PASP*, 114, 144
- Fruchter, A. S. 2011, *PASP*, 123, 497
- Gröchenig, K., & Strohmer, T. 2001, in *Nonuniform Sampling: Theory and Practice* (Springer)
- Hirata, C., & Seljak, U. 2003, *MNRAS*, 343, 459
- Takeda, H., Farsiu, S., Christou, J., & Milanfar, P. 2006, in *The Advanced Maui Optical and Space Surveillance Technologies Conference*, E27
- Werther, T. 1999, *Reconstruction from Irregular Samples with Improved Locality* (<http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.28.402>)