VLASS Memo 21: VLASS Cumulative Imaging

Juergen Ott

September 30, 2024

Abstract

This memo describes some basic image combination methods for data from the "VLA Sky Survey" (VLASS): linear mosaic and rms-weighted averaging in the image domain, as well as joint deconvolution of combined visibilities using different visibility gridders. All of the cumulative imaging methods reach lower rms values, close to what is expected, they also retain the fluxes and spectral indices of their input images. Whereas the image domain methods can be done with minimal computing effort, the images may retain some artefacts from the input images, and the average point spread function is generally (unless smoothed) not a Gaussian. Astrometric corrections have been applied to the input images and are inherited for the cumulative image. Joint deconvolution methods, on the other hand, are generally very slow as they require the 'aw-projection' gridder in CASA. Overall, they may improve the image fidelity and allow for deeper and more flexible imaging. We recommend weighing images based on $1/\sigma^2$ (σ the rms per pixel) and then averaging. This should be done for Single Epoch (both Taylor Terms) and Coarse Cube images by the observatory; Quick Look images can be combined on request. A python script for the combination is described here and has been provided to the VLASS team. With the advent of faster, GPU-based awp2 algorithms, joint deconvolution of the combined visibilities is recommended. Optimized cleaning parameters should be derived once the performance of awp2 is adequate.

1 Introduction

The VLA Sky Survey (VLASS; Lacy et al. PASP, 2020, 132, 1009) is conducted in three different epochs, each of them covering the visible sky from the VLA site. All epochs are observed in on-the-fly mode, which takes data while the array is scanning across a defined observing area, or 'tile'. This mode dumps a large number of individual pointings (defined by a sequence of phase centers) into the correlator that are then imaged together as a mosaic. After calibration, quick look (QL) images are obtained for $1^\circ \times 1^\circ$ areas. The data are later reprocessed into 'Single Epoch' (SE) images, including self-calibration, and improved imaging techniques. The SE images also apply a two-order Taylor Term expansion for the wideband images, which essentially provides a flux and

a spectral index map. In addition, 'Coarse Cube' (CC) images will be provided, one 128 MHz spectral window per plane, without Taylor Term expansion. Eventually the data from all multiple epochs will be combined into deeper images through cumulative imaging. A first sketch of cumulative imaging has been provided by Steven Myers[1]. Cumulative imaging shall provide deeper images of the sky, ideally with $1/\sqrt{n}$ times lower rms (where n is the number of epochs included if they have similar noise properties). Cumulative images can also serve as a reference image for transients.

2 Cumulative Imaging Methods

In this memo we apply several options to combine the data. The methods break down into two branches: Image plane combination and joint imaging of the combined visibilities.

2.1 Image Plane Combination

We test two image plane combination methods. Given the superiority of the SE images over QL, we use the former for our tests. The basic combination methods that we test here are linear mosaicking and rms weighted averaging. The main advantage of these methods over joint imaging is that they are computationally very fast. However, they do not take advantage of the improved, combined uv-coverage that is available to joint imaging, and artefacts from the individual SE images will be propagated into the final product. The combined point spread function (psf) is also less defined.

2.1.1 Linear Mosaic

The linear mosaic method consists of averaging pointings together, using the primary beam as a weighting function. This method is preferred when each mosaic point was imaged and cleaned independently. The linear mosaic combination can be written as follows:

$$I_{\rm comb} = \frac{\sum_i PB_i\, I_i}{\sum_i PB_i} \tag{1}$$

where I_i are the individual images, and PB_i the individual primary beams. Linear mosaicking is therefore a primary beam weighted averaging technique. In CASA's LINEARMOSAIC tool it is also possible to select different weighting functions depending on whether primary beam corrections to the input fields were applied or not. CASA also offers a 'Sault' weighting scheme to downweigh data at the edges of a mosaic. Note that for the tests in this memo, the images from each epoch were already mosaicked in a previous SE imaging run. The weighting function of adding epochs are therefore the mosaicked primary beams of each epoch; in other words, the sensitivity function across each input image.

[1] *https://docs.google.com/document/d/1t0QcH5Xd3bWwEqxV3DyR7N7FzncC7Z2mrcFm7l6RM_Y/edit*

2.1.2 RMS-Weighted Averaging

The sensitivity of each mosaicked SE image is also expressed in an rms (root mean squared) sensitivity map. This map was constructed by the CASA task IMDEV which determines the rms over a defined area around each pixel and assigns the pixel value to that number. The sensitivity map therefore not only reflects the depth and primary beam at each position, but also shows increased rms values where sidelobes and other artifacts still persist after deconvolution. The image combination uses the inverse of the squared rms map as a weighting function for averaging of the epochs: The rms weighted combination is done via:

$$I_{\rm comb} = \frac{\sum_i 1/\sigma^2 I_i}{\sum_i 1/\sigma^2} \tag{2}$$

(where σ is the rms per pixel) Regions of high noise in each map are therefore suppressed. Outlier rejection is possible (median in its simplest form), but there are limitations given the low number of input VLASS images (typically three).

2.1.3 Other Image Domain Combination Methods

Other image based combination methods that are not tested in this memo are, e.g., to combine the model data, convolve with a common beam, and adding in a weighted combination of the residuals. This method requires the model images from the cleaning process, which are not archived. However, models can also obtained through deconvolution or source finding algorithms. Also artificial intelligence techniques may increasingly be used to obtain the best guess image based on individual input images. But we consider any of these methods beyond the scope of this memo.

2.1.4 Point Spread Function

Each input image has its own point spread function (psf) and the image plane combination techniques do not offer an a priori new psf. One may assume though that averaging the psfs (e.g. weighed by $1/\mu^2$; where μ is the median across each input rms map) from the individual epochs would result in a psf that matches well the output image. Such an average, however, is not a 2D-Gaussian anymore. Users should be aware of this and use the combined, non-Gaussian psf image whenever possible for their analysis. Nevertheless, a Gaussian fit to the combined psf pattern should be added to the header as it is an approximation to the combined psf.

Of course it is always possible to smooth each input image to a beam size that matches the largest psf of the input images. This will provide a fairly clean psf, but the resolution of the output image will be degraded.

2.1.5 Spectral Index

VLASS SE images are calculated using Taylor Term expansion in the spectral domain with two Taylor terms. CASA's TCLEAN uses this mode for wideband

imaging to derive the spectral index α of maps via: $\alpha = \mathrm{tt1/tt0}$, where tt0 and tt1 are the 0th and first order Taylor term images, respectively. Any cumulative image combination technique is applicable to both Taylor terms, and α, derived from combined tt0 and combined tt1 images can therefore be calculated straightforwardly.

2.1.6 Astrometry

CASA's TCLEAN 'mosaic' gridder is used for the bulk of the SE imaging. This gridder, however, neglects the curvature of the sky, omitting calculations for the w-term of the Fourier transform of the visibility function. The 'mosaic' gridding method is fast, but shifts the position of each source in the map as a function of wavelength. For a weak w dependency, at smaller zenith angles, this effect produces an astrometric error and sources in the image are displaced. For larger zenith angles, a chromatic error will also deteriorate the spectral index of each source. VLASS memo 14 (M. Lacy)[2] presents a detailed analysis of the astrometric effects and an analytic correction function for the source displacement. This correction is applied to the bulk of VLASS fields as a lateral adjustment to the FITS image reference pixel. Regridding the images to match the world coordinate systems of the pixel grid is therefore needed before combination.

For the largest zenith angles, the chromatic errors that affect the spectral indices are more significant and TCLEAN's 'aw-projection' gridder is required to properly correct for the w-term of the sky curvature Fourier transform. As compared to the 'mosaic' algorithm, 'aw-projection' is very slow, and it would be prohibitive to image the entire sky with this gridder. An improved 'awp2' algorithm implemented on GPU-based computing infrastructure, however, is currently under development and may overcome the speed limitations in the future.

2.2 Joint Imaging

The joint imaging technique does not use previously obtained SE images, but uses the calibrated visibilities of each epoch together to create a new image. The advantage is that this method utilizes the combined uv-plane and thus produces a better defined psf and image fidelity. This aides the deconvolution process and should produce fewer artifacts. Astrometric errors for different hour angles and declinations, however, cannot be corrected individually anymore and the 'aw-projection' gridder is typically required for creating the images. The improved depth of the combined image also allows for deeper cleaning and clean boxes for fainter sources. In addition, joint imaging opens up the possibility to optimize the imaging parameters to better suit the combined images, including visibility weighting, the application of multi-scale or ASP algorithms, etc. All of that should combine to a much better image product than image plane combined SE images. Variable sources, however, cannot be deconvolved without errors, when their visibilities are imaged together.

[2] https://library.nrao.edu/public/memos/vla/vlass/VLASS_014.pdf

Given the need for aw-projection, the cost for the re-computation of the combined images is much higher and can be up to several weeks with the current implementation of 'aw-projection' and 32 w-projection planes on 16 CPU cores (see below). As mentioned earlier, GPU-based 'awp2' gridding may reduce this to a level that can be used in production. The joint imaging, however, will always be slower than the image based combination.

For this memo, we have used three different joint imaging gridders for VLASS: the 'mosaic' gridder, 'aw-projection' with a single projection plane, and 'aw-projection' with 32 projection planes.

2.2.1 Mosaic Gridder

The 'mosaic' gridder combines all the pointings in a single uv-plane, applying the respective phase gradients to each offset pointing and a prolate spheroidal as the gridding kernel. The combined uv-plane is then inverted and can be primary beam corrected where the primary beam is also the combination of the primary beams of each pointing, adjusting for different depths of observations. It is therefore more like a theoretical sensitivity map. The mosaic gridder does not correct for w-terms and thus introduces some astrometric errors at large zenith angles. These data will need astrometric correction, potentially based on the average corrections for the individual epochs. Given the different zenith angles of the input data, however, there is no straightforward solution to this problem.

2.2.2 aw-projection with a Single Projection Plane

aw-projection is a gridder in CASA's TCLEAN that takes into account the w-term for non-planar sky projection as well as the illumination patterns of the dish that is restricted by the feed legs and other telescope features. The w-projection planes represent the number of discrete w-values to be used to quantize the range of w-values present in the dataset. The number N_w of w-projection planes that are needed can be calculated as $N_w = 0.5 \times W_{\rm max}/\lambda \times l$ where $W_{\rm max}$ is the maximum w in the data, λ the wavelength, and l is the image size in radians; cf. CASA docs[3]). For a full treatment, 32 w-projection planes are adequate for VLASS imaging. The aw-projection with a single plane should be taken like a test of the underlying gridding method, without an actual correction, similar to the 'mosaic' gridder.

2.2.3 aw-projection with 32 Projection Planes

With 32 w-projection planes, the VLASS data should be adequately sampled in the w-plane and w-term errors of the astrometry and chromatic spectral index issues should be eliminated. This method, in principle, should produce the

[3]https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html#Wide-Field-Imaging

best joint image, with a newly computed clean beam. The calculation of the aw-terms, however, is very computationally expensive at this time.

3 Test Fields, Data Preparation, and Image Combination

We study two fields to determine the best cumulative imaging technique, a GOODS-North (GOODSN) field centered at $\alpha = 12h41m3.7s$, $\delta = +64°30'0.0$" (J2000) and one toward the Chandra Deep Field South (CDFS) at $\alpha = 3h32m28.0s$, $\delta = -27°48'30.0''$ (J2000). GOODSN is a fairly northern field where astrometric errors will be small. The southern CDFS, however, should expose any astrometic errors more prominently, as they depend on the zenith angle of the observations as taken with the northern hemisphere VLA.

To prepare the data for the tests, we split out the visibilities $1° \times 1°$ around each central coordinate followed by applying the SE imaging script to the individual data sets. After a primary beam correction, the images are then corrected for astrometric errors, according to VLASS memo 14. This produces the individual images for each observation. Note, that VLASS production only archives a selected number or imaging products. In particular, only primary beam corrected images and rms images are stored in the archive; the primary beams, dirty psf, non-primary beam corrected data, etc. are not available without re-imaging.

For joint imaging, we are not using the CASA pipeline implementation of the imaging script, but the script that was derived from the 'VLASS Imaging Project' (VIP), a script that was later adapted for the SE production pipeline. We modified the script, however, such that each individual observation was self-calibrated individually (using the joint images as model). The STATWT step to recalculate the weights is also applied for each dataset individually. All visibilities were then jointly imaged in TCLEAN. Note that due to aliasing, the aw-projection images had to increase in size from 12500x12500 to 16384x16384 pixels.

4 Results

4.1 Imaging Results and Quality Comparison

VLASS images are notoriously large with small sources in them. So rather than showing the full images, we show cut-outs around selected sources for each field. For GOODSN, tt0 images are shown in Figs. 1-6. tt1 images, corresponding to the sources depicted in Figs. 1 and 2 are shown in Figs. 7 and 8. The CDFS cutouts are only shown for the tt0 component in Figs 9-14. All panels show the individual images for each epoch (labeled E1, E2, and E3), as well as the results from the various image combination methods (lin-mosaic: linear mosaicking, w-average: $1/\sigma^2$ weighted average, j-mosaic: joint imaging mosaic gridder, j-awp1:

joint imaging aw-projection gridder with 1 w-projection plane, j-awp32: joint imaging aw-projection gridder with 32 w-projection planes). Each figure has all its panels on the same scale.

Inspection of the images shows that all image combination methods improved the image quality. It turns out that the rms weighted average is one with the lowest image artefacts. This may be understood given that noisy and sidelobe-corrupted areas are weighted down. For the joint imaging, the mosaic gridder shows the best image quality.

4.2 RMS Measurements

The rms values of the individual and the combined images are shown in Table 1, together with the respective beam sizes. Since the images from joint mosaicking have a newly calculated beam,we also show the rms values after smoothing to the psf derived from the average of the individual images, the psf that is derived for the image plane combination. rms values are calculated from areas $\alpha = 12h45m10.8s...12h40m04.4s, \delta = 64°29'02''...64°51'06.0''$ for GOODSN and $\alpha = 3h36m09.8s...3h34m12.0s, \delta = -28°14'58''...-28°01'21''$ for CDFS, regions chosen to have virtually no detectable sources in them, away from the image edges.

The linear mosaicking technique shows slightly better rms values for GOODSN (tt0 and tt1), but is inferior for CDFS. As mentioned above, however, the rms weighted average appears to have somewhat better image fidelity and is therefore the better technique for image plane combination. Furthermore, the rms weighted average is remarkably close to the expected, combined rms from the individual images.

Joint imaging shows that the mosaic gridder generally produces better rms values than the aw-projection gridder. It is slightly above the theoretically combined rms but also shows somewhat higher values than the rms weighted average. Smoothing to the combined, averaged, clean psf used for the image domain, only changes this behavior slightly.

We should note though that given the extreme computation time, in particular for aw-projection, a full exploration and optimization of imaging parameters for joint imaging techniques were not undertaken for this memo. Also, 'awp2' may change some of the image parameters since the algorithm has been modified, in particular using cube-based calculation of the spectral index rather than Taylor Term expansion. It is the goal of a future memo to derive the best parameters for the joint visibility imaging.

4.3 Flux Densities

In Figs. 15-18 we compare the flux densities of individual sources in each combined image with those of the individual input images. Overall, we do not see any differences or systematic trends between the imaging techniques, all fluxes are well represented in the combined images. Some scatter certainly remains, but appears to be of statistical nature for both, tt0 and tt1 images.

	GOODSN			CDFS		
	rms (μJy/beam)		beam bmaj×bmin; PA	rms (μJy/beam)		beam bmaj×bmin; PA
	tt0	tt1		tt0	tt1	
image 1	124	639	3.08" × 2.22";−67.7°	148	772	4.99" × 2.26";−4.3°
image 2	125	648	2.69" × 2.12";26.0°	175	948	4.89" × 2.09";−9.2°
image 3	126	639	2.76" × 2.23";9.7°	150	768	5.08" × 2.18";−11.6°
expected, combined rms	72	371		90	472	
	average clean psf					
linearmosaic	70	361	2.54" × 2.45";4.1°	103	543	4.97" × 2.18";171.5°
weighted average	72	371	2.54" × 2.45";4.1°	91	474	4.97" × 2.18";171.5°
	joint imaging					
joint-mosaic	75	381	2.53" × 2.43";13.3°	95	497	4.95" × 2.19";−8.5°
joint-awp1	81	459	2.35" × 2.26";17.3°	118	695	4.58" × 2.02";−8.4°
joint-awp32	81	459	2.35" × 2.25";16.1°	118	694	4.58" × 2.02";−8.4°
	joint imaging smoothed to image plane clean psf					
joint-mosaic	75	383	2.54" × 2.45";4.1°	95	497	4.97" × 2.18";171.5°
joint-awp1	81	453	2.54" × 2.45";4.1°	121	689	4.97" × 2.18";171.5°
joint-awp32	81	453	2.54" × 2.45";4.1°	120	688	4.97" × 2.18";171.5°

Table 1: GOODSN and CDFS rms values and beam sizes for the individual images as well as for cumulative images, obtained through different combination techniques.

4.4 Astrometry

Astrometric corrections are dependent zenith angle, i.e. Hour Angle and Declination. Thus, the effect for CDFS should be much stronger than for GOODSN. And, in fact, the offsets calculated by the equations given in VLASS memo 14, show corrections $< 0.2''$ for GOODSN, which is less than about 10% of the beam. For CDFS, the corrections are of order 0.5" in declination, close to the pixel size of 0.6" and about a quarter of the minor axis beam size. As shown in Table 2, the image domain combination indeed eliminate most of the astrometric offset (in the table the reference source has been taken from NED). The joint imaging techniques, indeed do not show that correction in δ. Surprisingly, the aw-projection with 32 w-projection planes did not resolve the astrometric offset entirely. It is, however, more accurate than a single projection plane, or the mosaic gridder. It is possible that more w-projection planes will provide better astrometric accuracy. Applying the mean corrections for the input images, however, should provide an acceptable astrometic accuracy given the beam sizes of the observations.

Image	$\Delta\alpha$ (arcsec)	$\Delta\delta$ (arcsec)
Image 1 (uncorr)	-0.10	-0.44
Image 2 (uncorr)	-0.39	-0.14
Image 3 (uncorr)	-0.15	-0.57
Image 1 (corr)	-0.21	0.04
Image 2 (corr)	-0.50	0.34
Image 3 (corr)	-0.30	-0.07
linearmosaic (corr)	-0.33	0.09
rms weighted average (corr)	-0.30	0.06
joint-mosaic	-0.15	-0.45
joint-awp1	-0.15	-0.50
joint-awp32	-0.18	-0.29

Table 2: Offsets of fitted Gaussians to the source SWIRE J033051.41-273013.7 in arcseconds. Fitting errors are typically 0.005" in α and 0.1" in δ. Beams are shown in Table 1. Note that the NED reference is not necessarily centered on the radio point source. Note that average offset corrections are $\Delta\alpha = -0.10$", $\Delta\delta = 0.483$" for image 1, $\Delta\alpha = -0.11$", $\Delta\delta = 0.484$" for image 2, and $\Delta\alpha = -0.15$", $\Delta\delta = 0.498$" for image 3 (offsets for GOODSN are calculated to be < 0.2" and thus less than about 10% of the beam).

5 Conclusions

In this VLASS memo, we discuss and test different image combination methods for cumulative imaging. Our tests include the image domain combination methods linear mosaic and rms weighted averaging, as well as joint imaging techniques using the mosaic gridder and aw-projection with 1 and 32 w-projection planes. Our results are as follows:

1. The image domain combination techniques are very fast compared to joint imaging. The differences span a few seconds for weighted averaging, minutes for linear mosaics, days for the joint imaging mosaic gridder, and a month for aw-projection with 32 w-projection planes.

2. The image quality overall is good. Slightly better images are provided by the weighted averaging techniques for the image domain, and using the mosaic gridder for the joint imaging technique.

3. This is also reflected in the rms, where the rms weighted average is very close to the theoretical rms as calculated from the input images. Also the mosaic gridder shows the lowest rms values for the joint imaging techniques.

4. Differences in flux densities between the input images and any combined image are comparable and mostly statistical nature.

5. The image plane methods naturally inherit the astrometric corrections from the input images. For joint imaging, the 'mosaic' gridder cannot account for astrometric errors, given the different zenith angles of each observation. In some cases, an average correction may be calculated from the input images, if the error would be a small fraction within the beam size. The aw-projection with 32 w-projection planes was not able to remove the astrometric inaccuracies entirely and additional w-projection planes may potentially be needed, increasing the computing time substantially over the already long processing times to many months per image as of the current NRAO DSOC cluster system, and without the 'awp2' GPU implementation.

6. The image plane combination has no Gaussian psf. An approximation by averaging the clean beams of each input image and re-fitting a Gaussian will likely be sufficient for most analysis needs, but the averaged psf itself should be provided for more accurate data analysis.

7. Spectral indices should be retained in the image combination by combining tt0 and tt1 images separately, and calculating $\alpha = \mathrm{tt1}/\mathrm{tt0}$.

6 Recommendations for VLASS

Based on the analysis presented in this memo, recommendations for cumulative imaging are as follows:

1. Provide a first cumulative image product using the $1/\sigma^2$ weighted average image plane combination technique. The combination is fast and the astrometry has already been taken into account for the individual images. Archival products will be sufficient as inputs for this technique. A caveat is that the psf is not Gaussian. However, the Gaussian fit to the average of the input clean beams should be added to the header and the accurate, combined psf (with $1/\mu^2$ median rms map weighting) should be provided. As an alternative, a combined image, smoothed to the largest input beam size may also be delivered, which has a lower resolution than the unsmoothed image, but it has better defined Gaussian psf.

2. When the 'awp2' algorithm is validated and when GPU-based computing facilities are available for a significant performance improvement, we recommend joint visibility imaging with the 'awp2' gridder. Imaging parameters should be optimized over the current SE pipeline.

7 Appendix:

A first data combination script in the image domain has been delivered to the VLASS project. The script follows the rms weighted averaging approach

and produces combined images that are a) the combination of the individual images, with a new, combined psf, and b) images that were first smoothed to the common largest beam and then combined. The script is run in CASA, and has been developed on CASA 6.5.4.9, which contains ASTROPY for plotting. The individual steps in the script are:

1. Reading in the images, tt0 and tt1, including the rms images - all of these are VLASS products that are available from the archive. If the rms maps are not available, the script will create them the same way as the VLASS Single Epoch (SE) pipeline.

2. As part of the SE pipeline, each SE image has an astrometric correction applied following VLASS memo 14. As a result, the input images for the data combination will not exactly align on a pixel by pixel basis anymore. Therefore, all input images are regridded to the first image in the list.

3. The square inverse of the rms map is used as the pixel-based weight for the rms weighted average of the input images. tt0 and tt1 images are combined separately.

4. The (clean) psf properties are read from the headers of the input images. Small images of the psfs are created and averaged together using $1/\mu^2$ (μ=median across each rms map) weighting. This combined psf is not a 2d Gaussian anymore, but will be approximated by a fitted Gaussian, the values of which are stored to the headers of the output images. The combined psf is also a product of the script and can be used for more accurate photometry if needed.

5. To avoid the uncertainties of a non-Gaussian beam, the script also first determines the largest major axis of all input image psfs. It then smooths all input images to a round psf with that largest axis, slightly extended by 5% to allow the smoothing algorithm to work properly.

6. The smoothed input images (tt0, tt1, as well as their rms counterparts), are now used for a $1/\sigma^2$ weighted average as above. The output image will have a better defined, although larger beam, as compared to the method that does not smooth the input images.

7. The script now creates new rms maps for all combined images, tt0, tt1, smoothed and unsmoothed.

8. For an evaluation of the final products, the script creates various diagnostic plots and statistics: of the input images (including rms image), output images, beam images, as well as difference images (combined images - each input image) and spectral index images. The plots together with image statistics are bundled on a webpage.

In the future, this script can be adjusted to:

- enable other image combination algorithms, e.g. with outlier rejection, or employing deconvolution algorithms.
- The script will also need to be adapted for non-SE images, such as Coarse Cube products.

Combinations in the visibility domain will be discussed in a separate VLASS memo.

8 Acknowledgements

I would like to thank Mark Lacy, Amy Kimball, Steve Myers, and John Tobin for discussions and critical reading of the manuscript and code review.

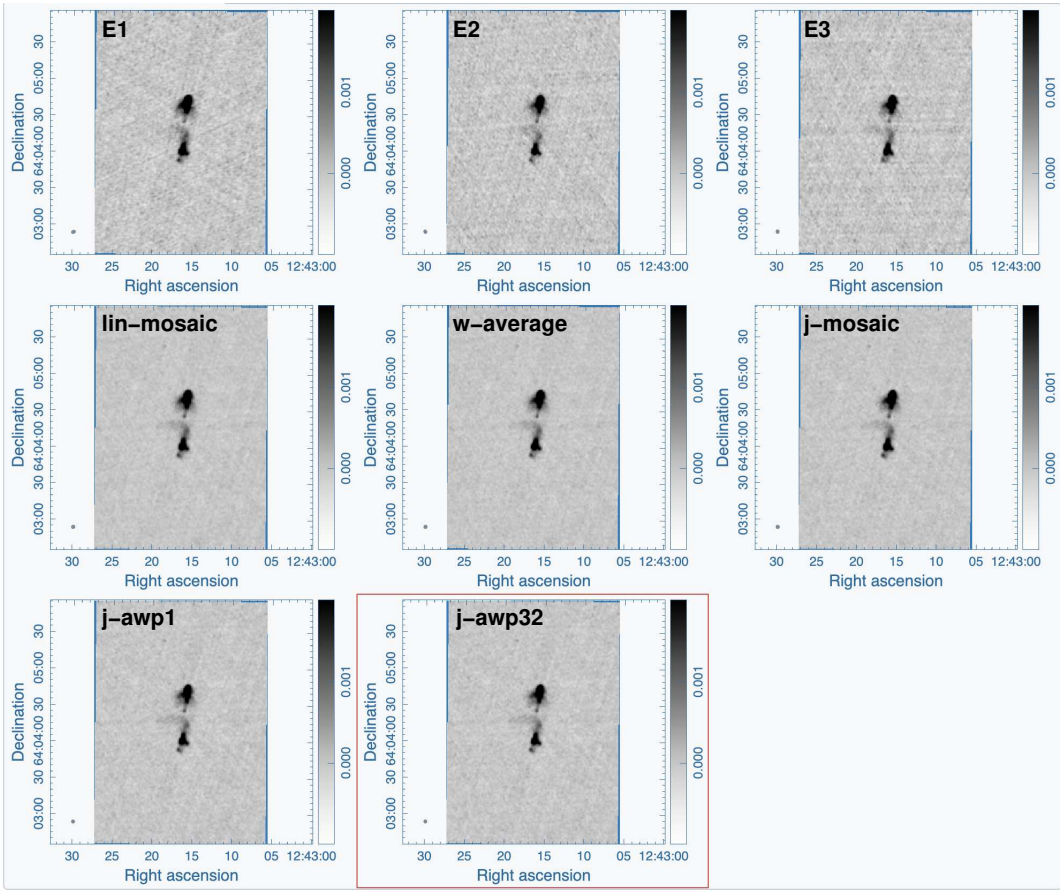

Figure 1: GOODSN, near source SDSS 6C B124107.3+642051: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

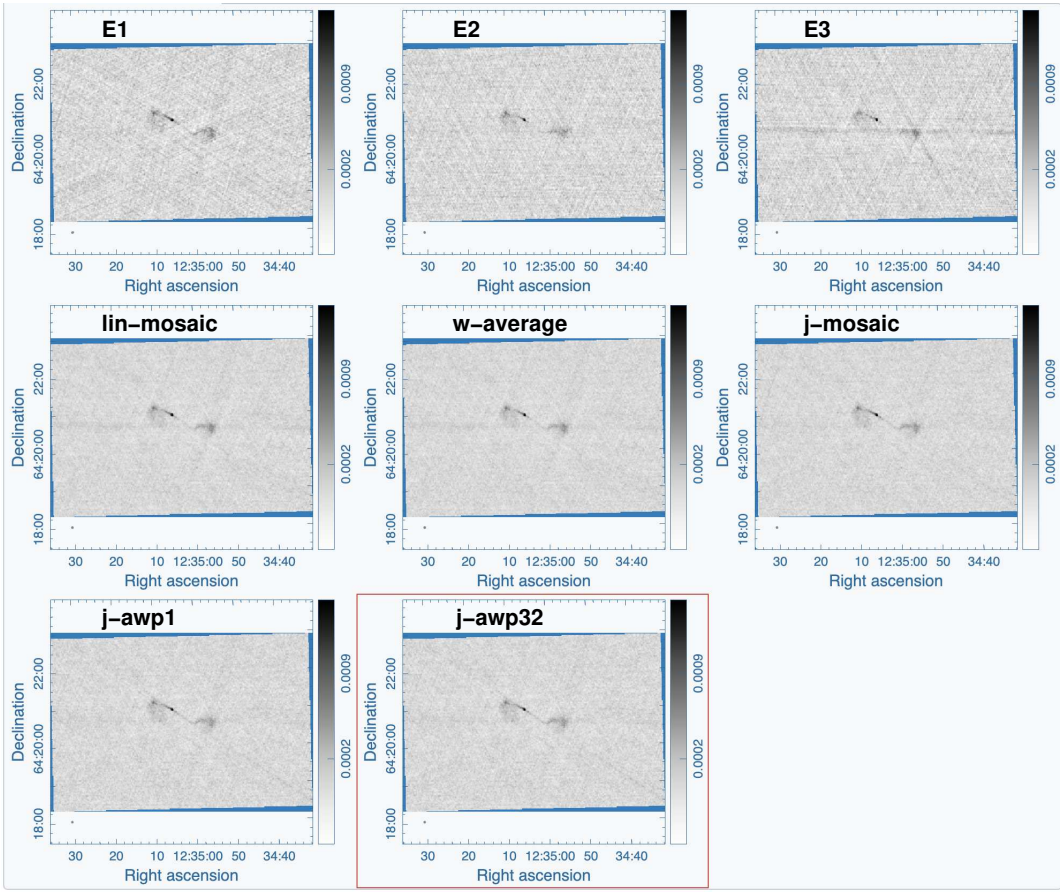

Figure 2: GOODSN, near source SDSS J123503.50+642105.2: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

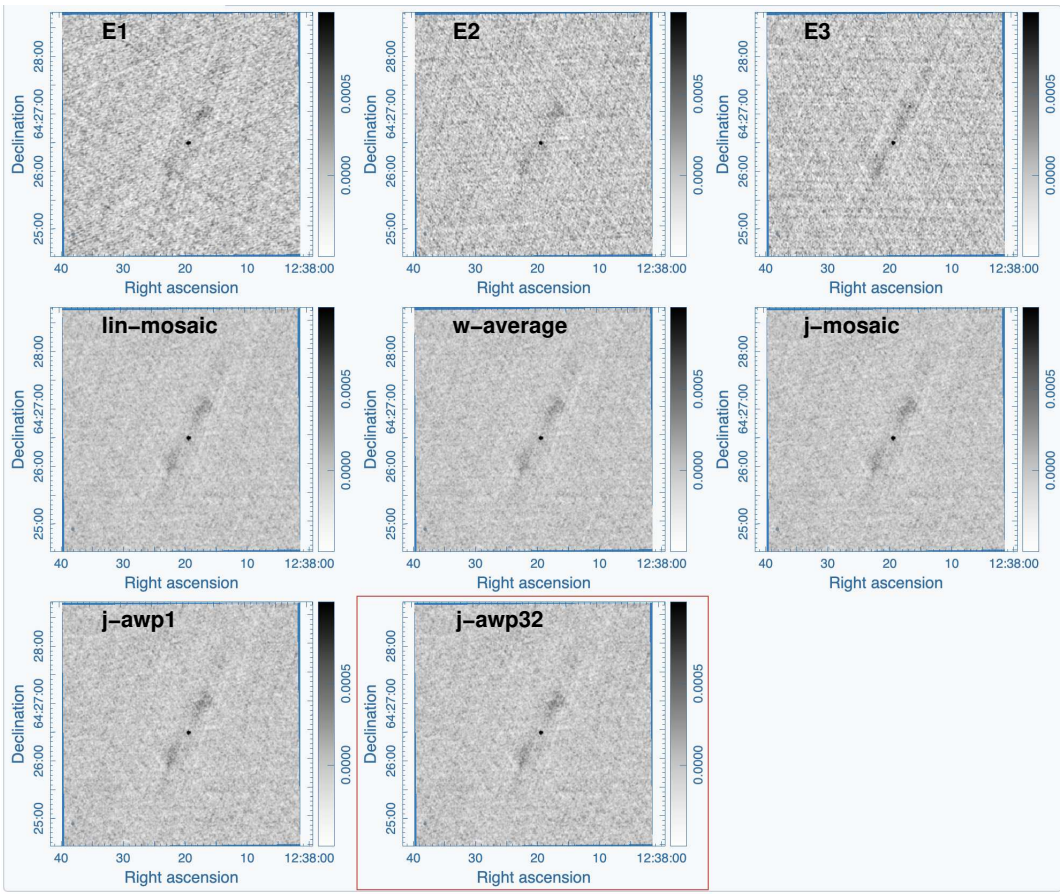

Figure 3: GOODSN, near source NVSS J123819+642630: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

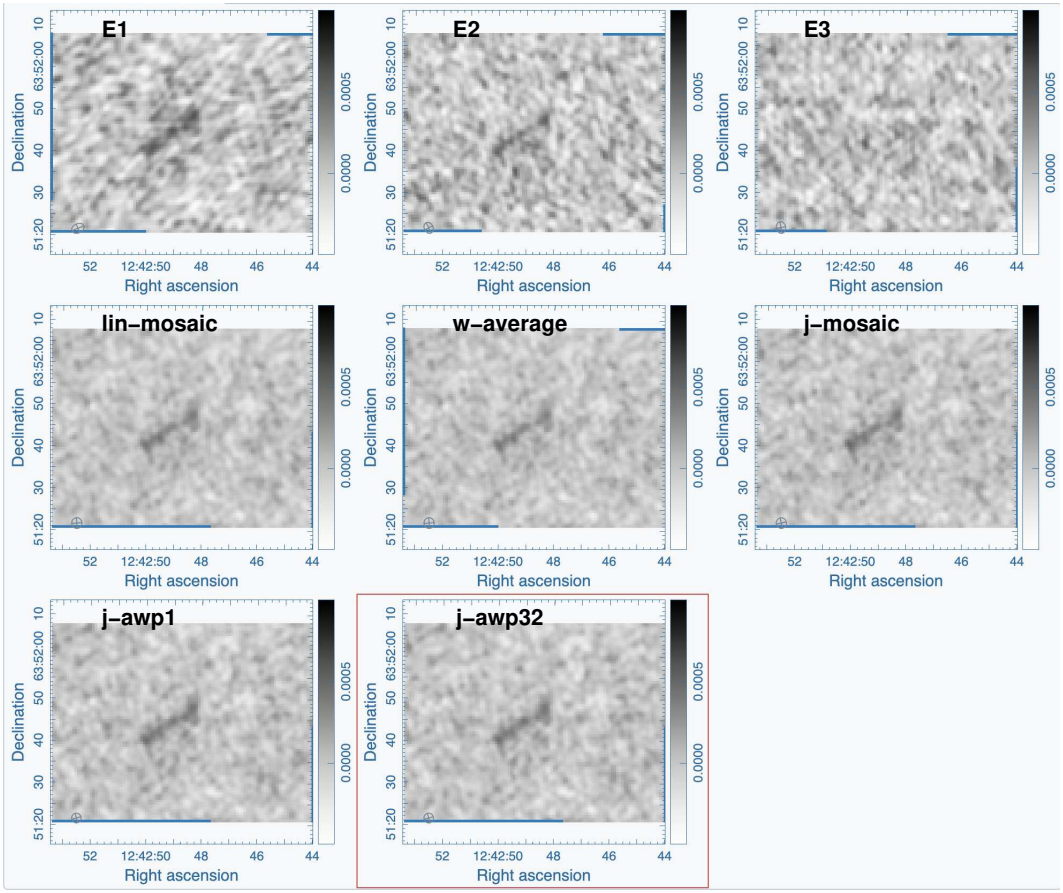

Figure 4: GOODSN, near source NVSS J124250+635140: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

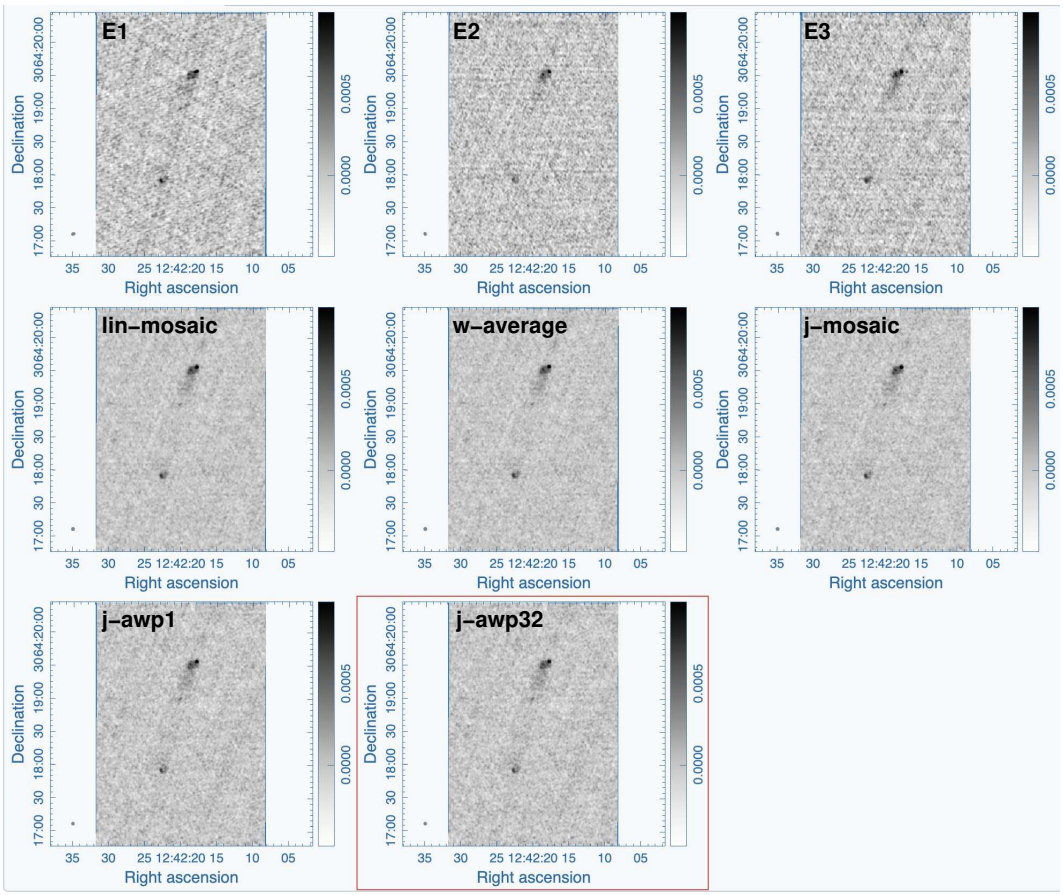

Figure 5: GOODSN, near source NVSS J124218+641926: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

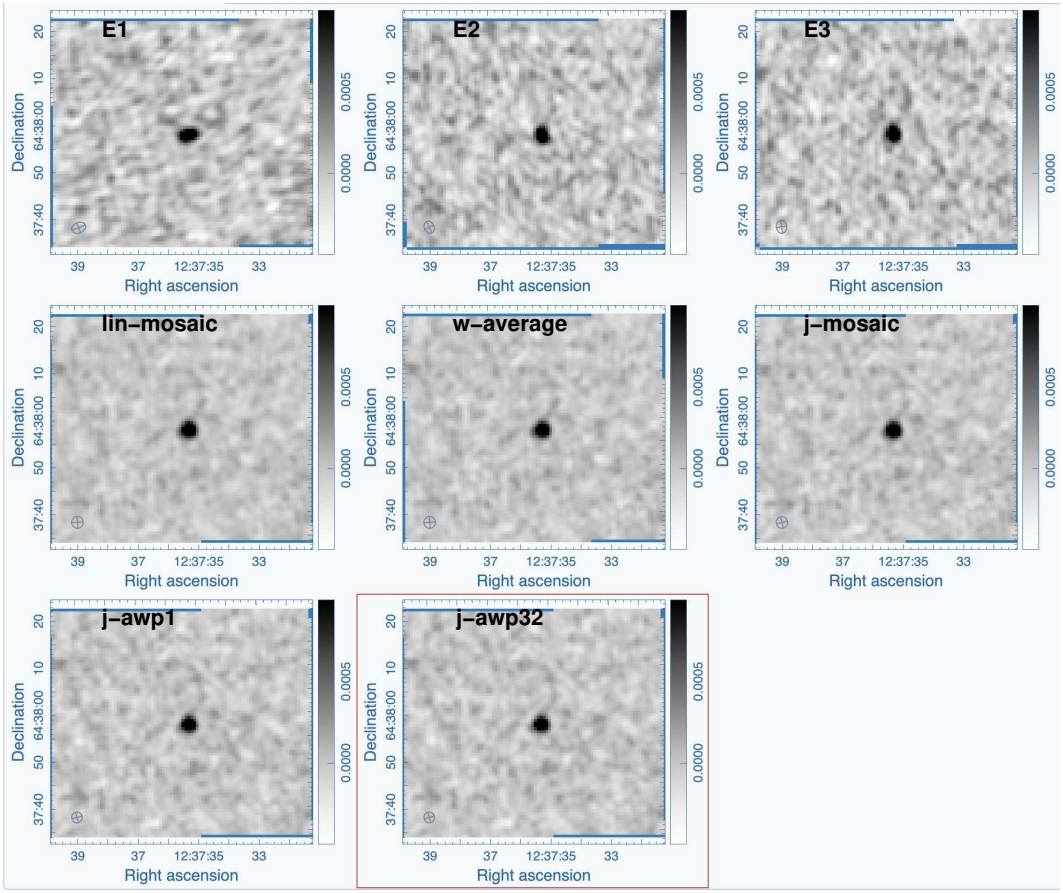

Figure 6: GOODSN, near source NVSS J123734+643757: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, *w*-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

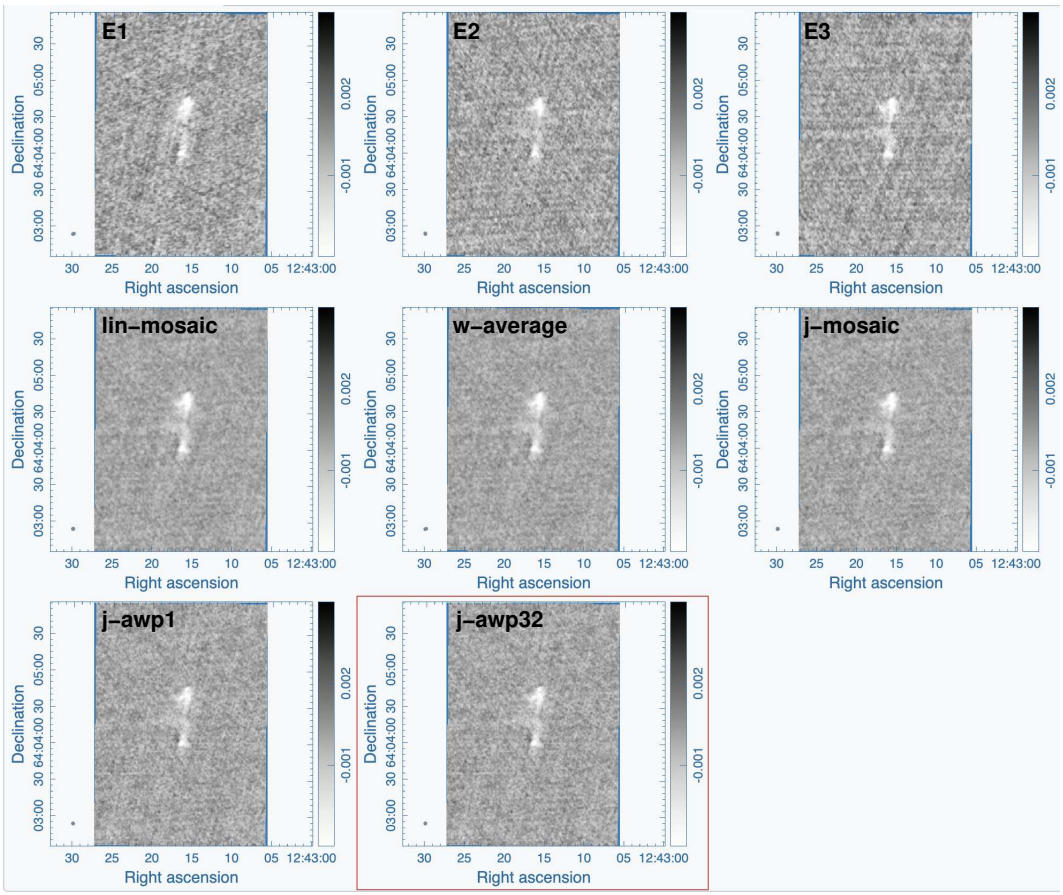

Figure 7: GOODSN, near source 6C B124107.3+642051: tt1 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

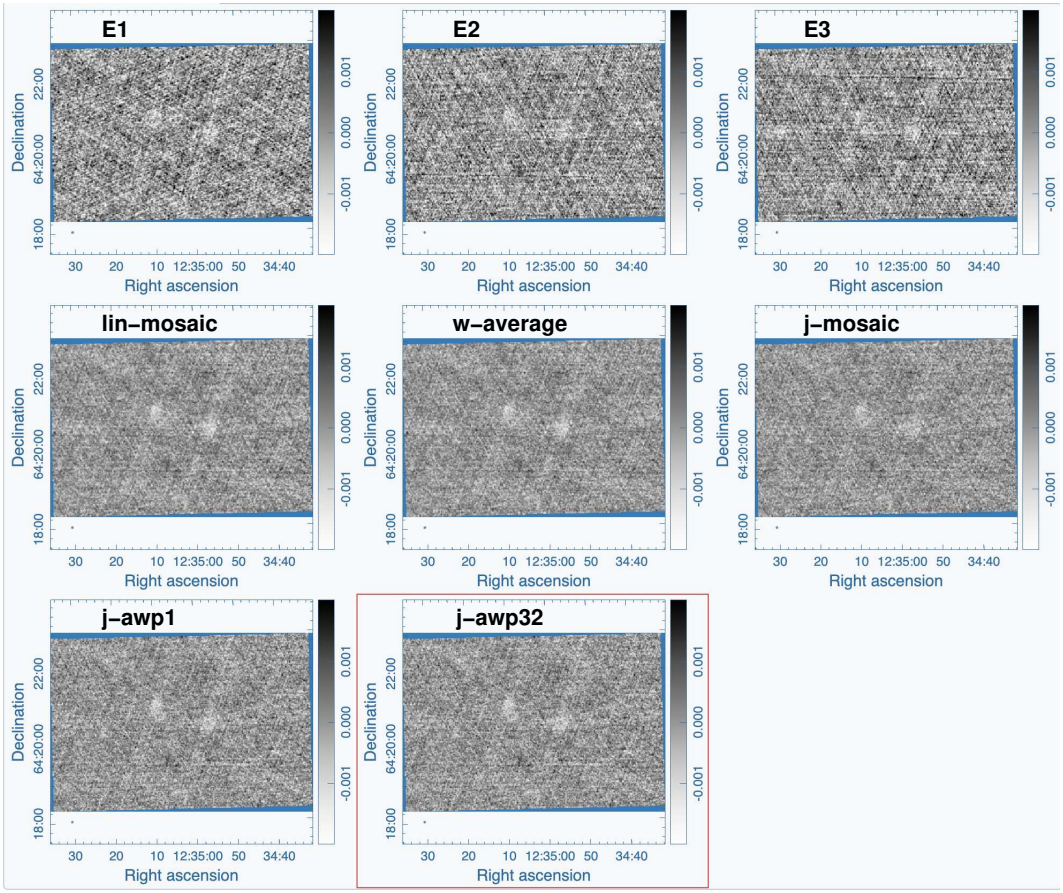

Figure 8: GOODSN, near source SDSS J123503.50+642105.2: tt1 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, *w*-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

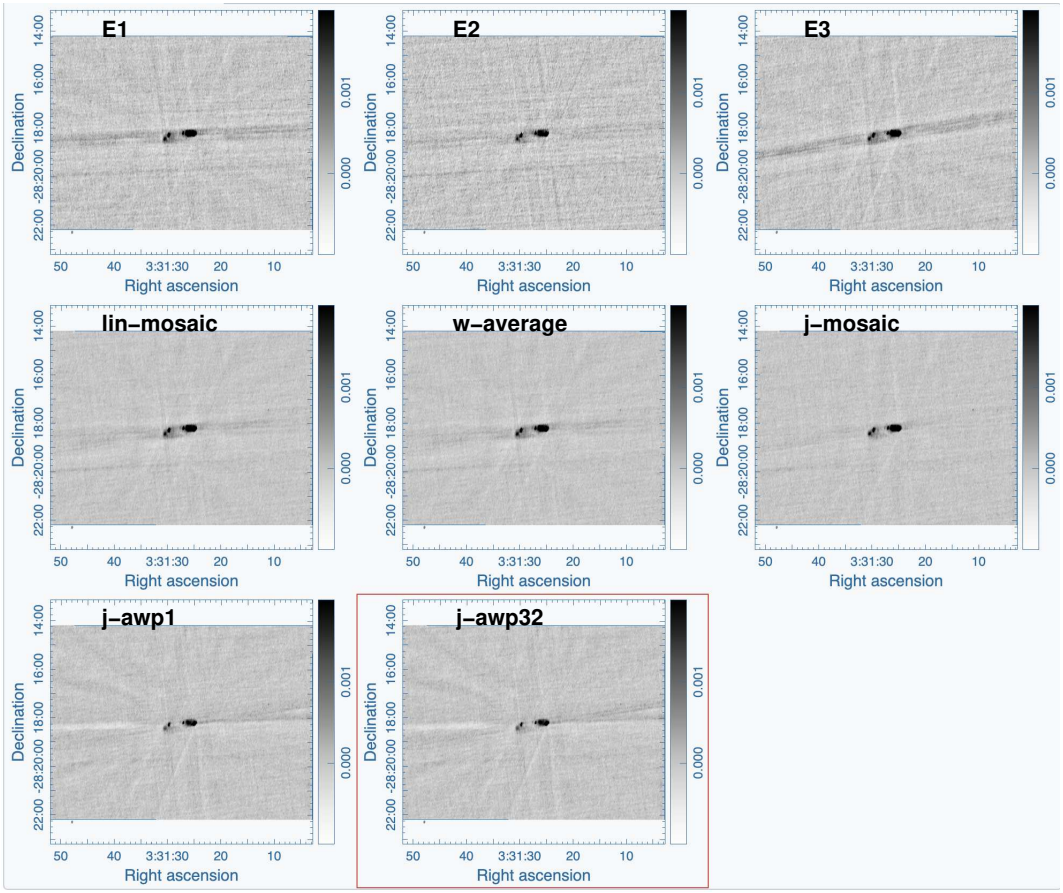

Figure 9: CDFS, near source AT20GDP J033128-281816: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

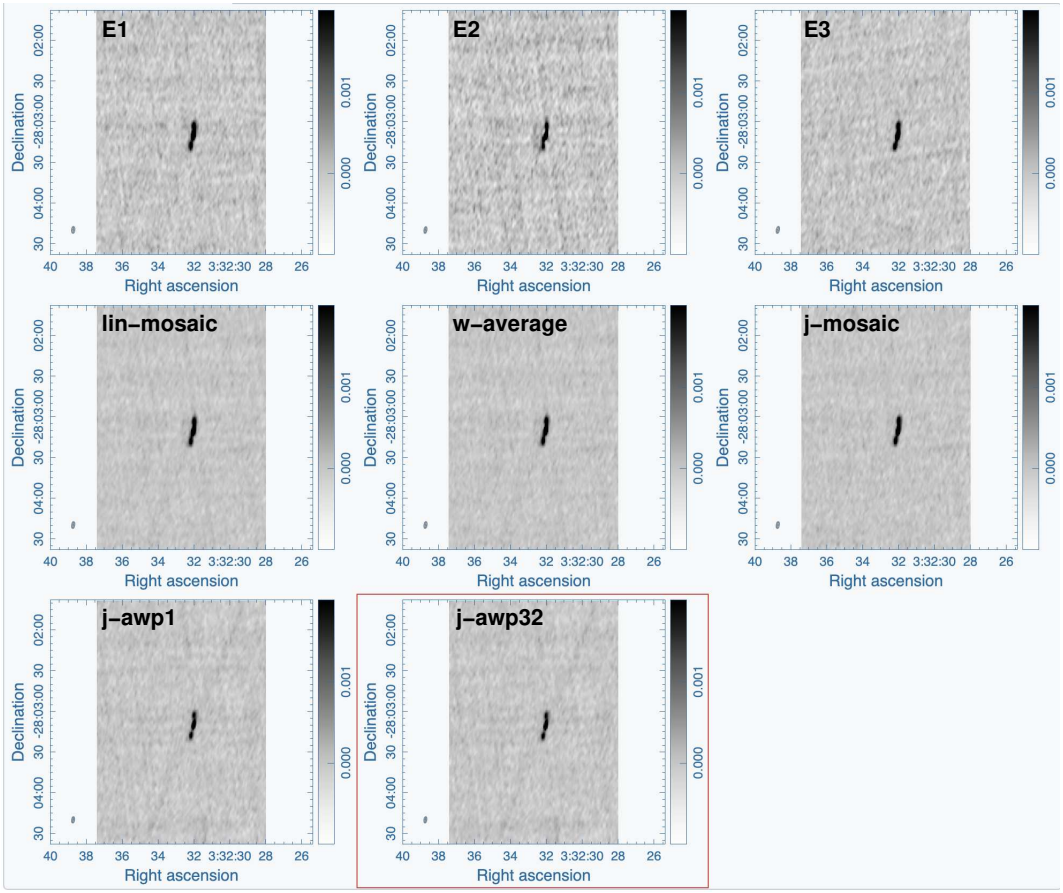

Figure 10: CDFS, near source ATLAS3 J033232.0-280305C: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

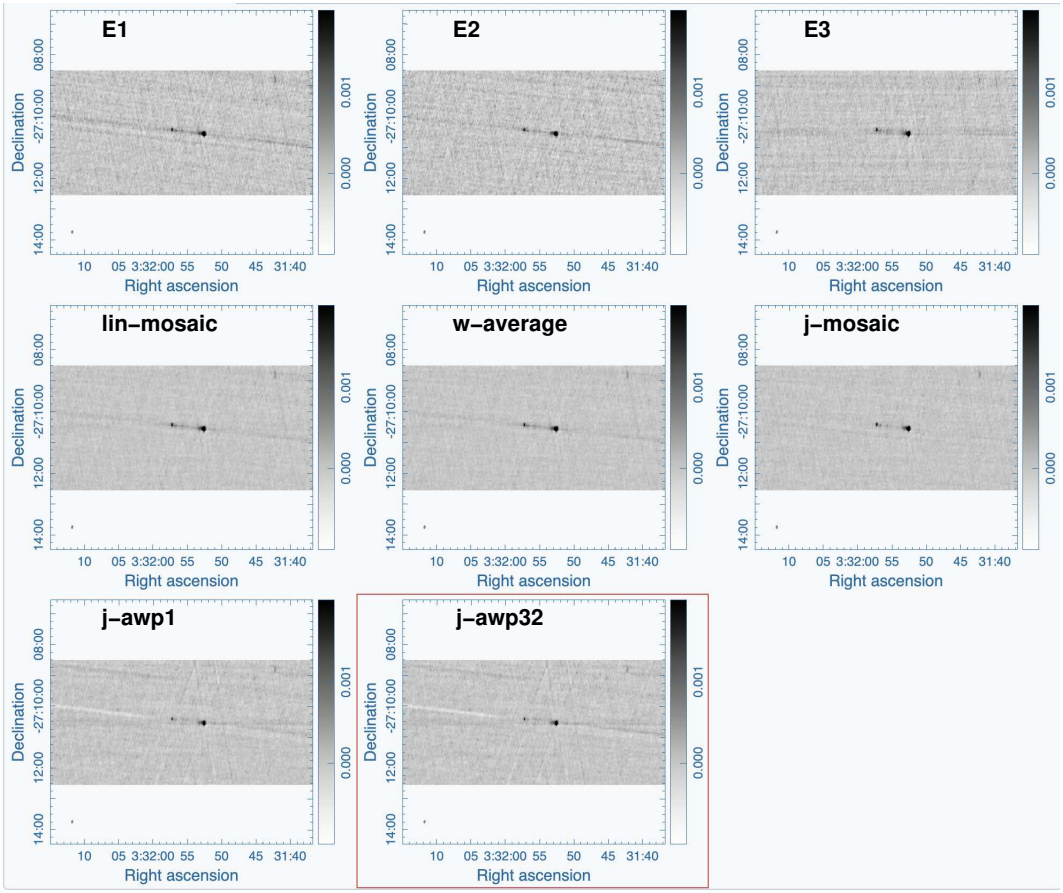

Figure 11: CDFS, near source VLSS J0331.9-2710: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

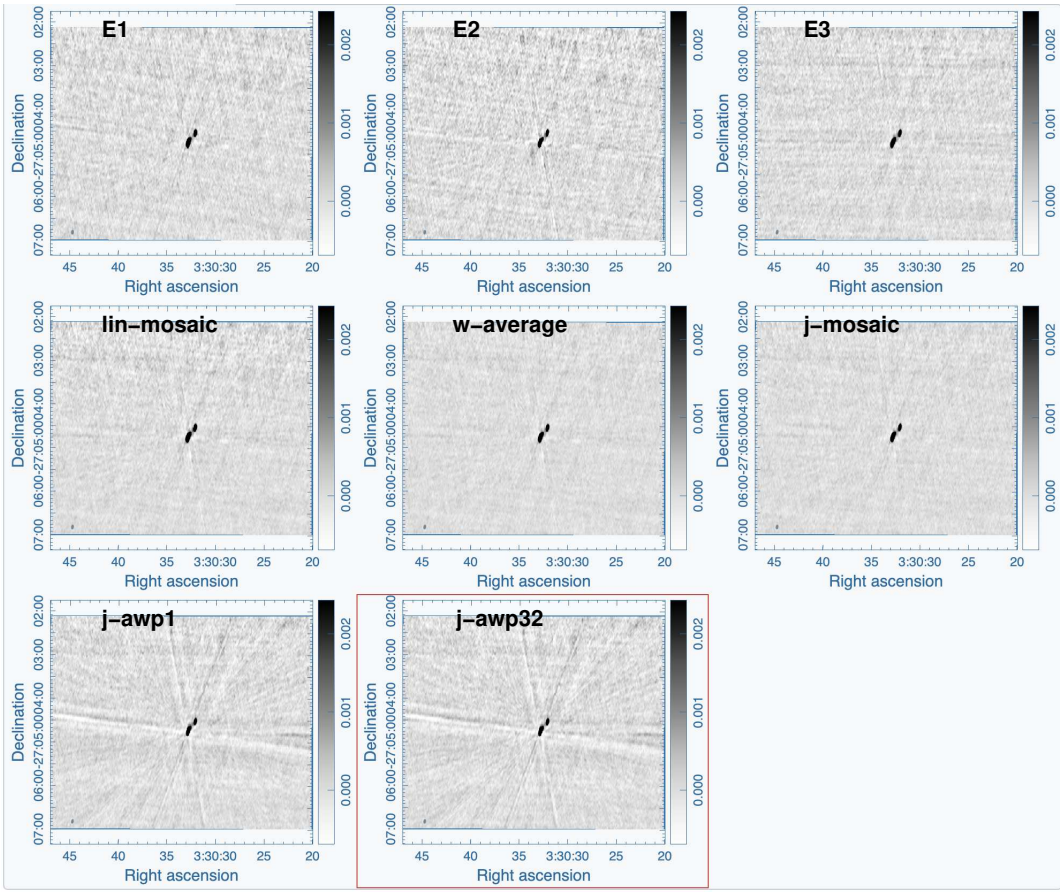

Figure 12: CSFS, near source PKS 0328-272: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

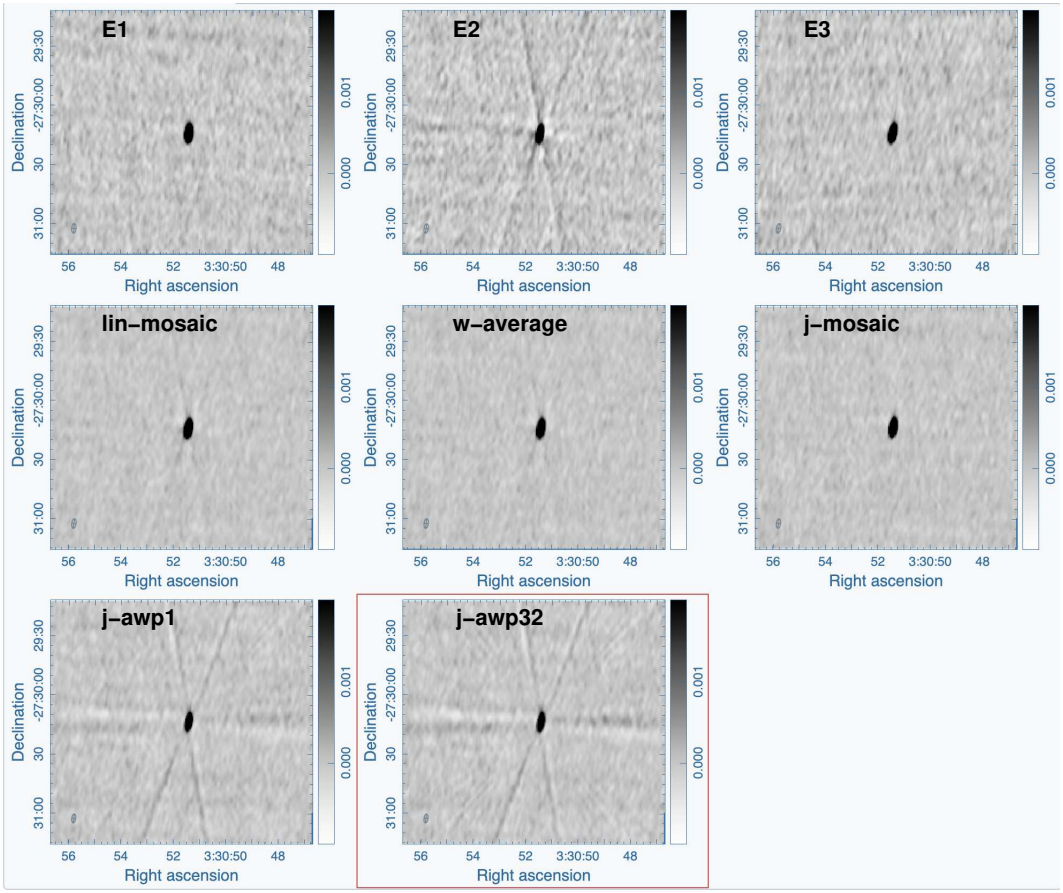

Figure 13: CDFS, near source SWIRE J033051.41-273013.7: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, w-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

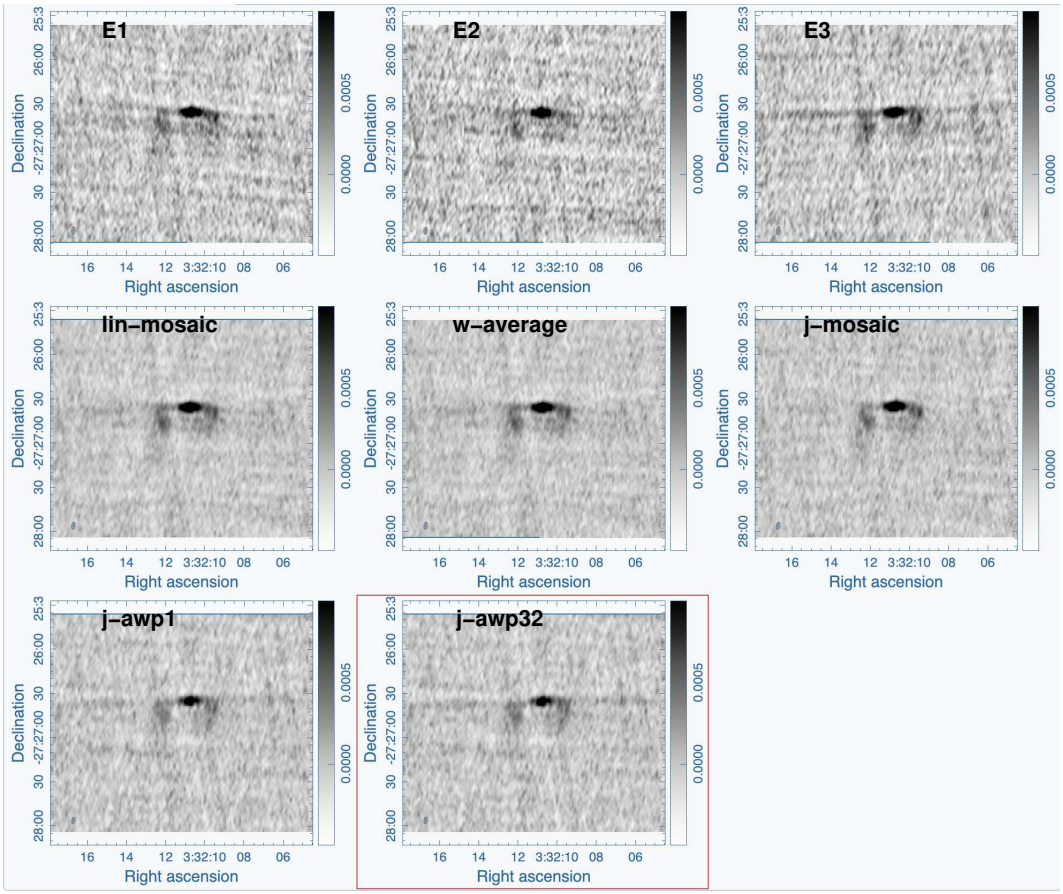

Figure 14: CDFS, near source NVSS J033210-272642: tt0 image results from the different imaging techniques. Top row are the individual epochs to be merged (E1, E2, E3). Merging results: lin-mosaic: image plane linear mosaic, *w*-average: image plane rms weighted average, j-mosaic: joint image mosaic gridder, j-awp1: joint image aw-projection gridder with 1 w-projection plane, and j-awp32: joint image aw-projection gridder with 32 w-projection planes. All images are on the same color scale.

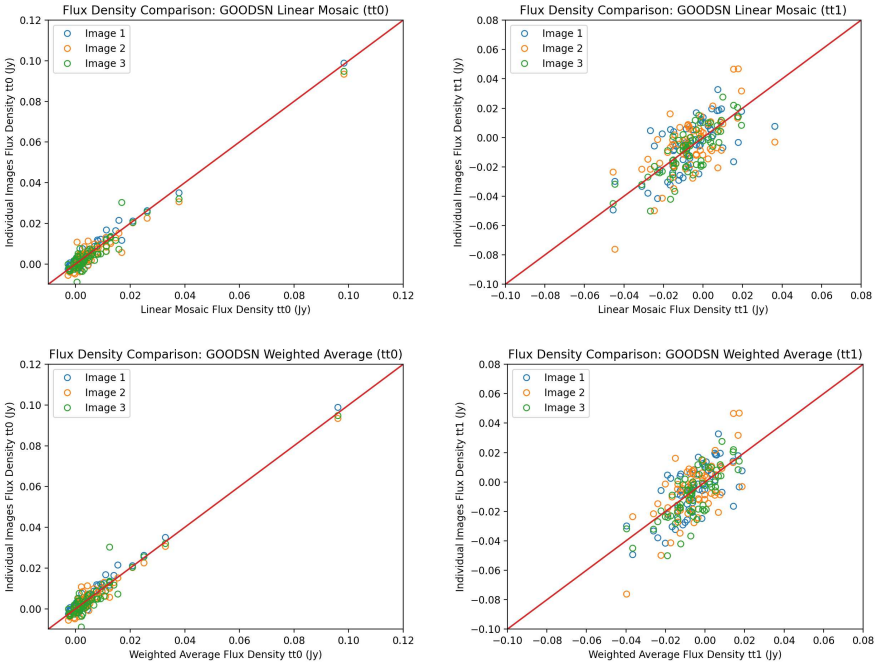

Figure 15: Flux density comparisons of the joint imaging combinations results versus the individual input images for GOODSN. **Top:** Linear mosaic, **Bottom:** rms weighted average. The left and right columns are the tt0 and tt1 image products respectively. The line is flux equality and not a fit to the data.

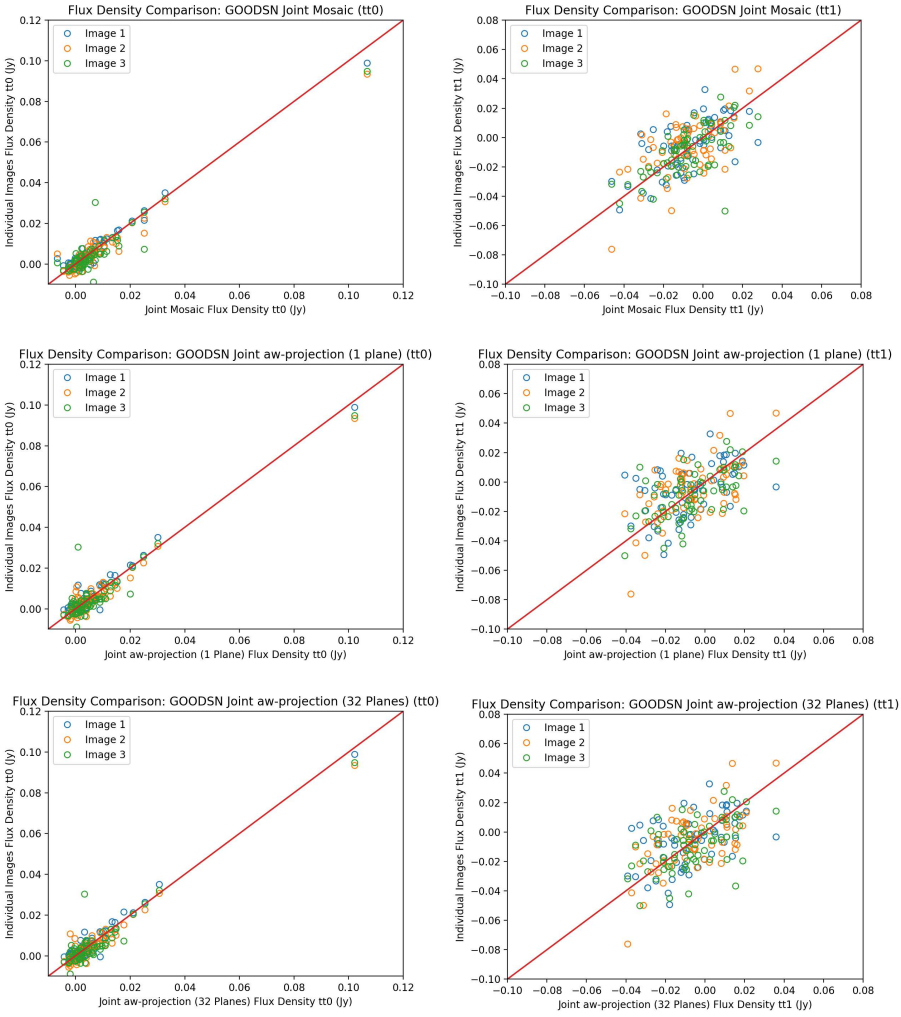

Figure 16: Flux density comparisons of the joint imaging combinations results versus the individual input images for GOODSN. **Top:** Mosaic gridder, **Middle:** aw-projection with 1 w-projection plane **Bottom:** aw-projection with 32 w-projection planes. The left and right columns are the tt0 and tt1 image products respectively. The line is flux equality and not a fit to the data.

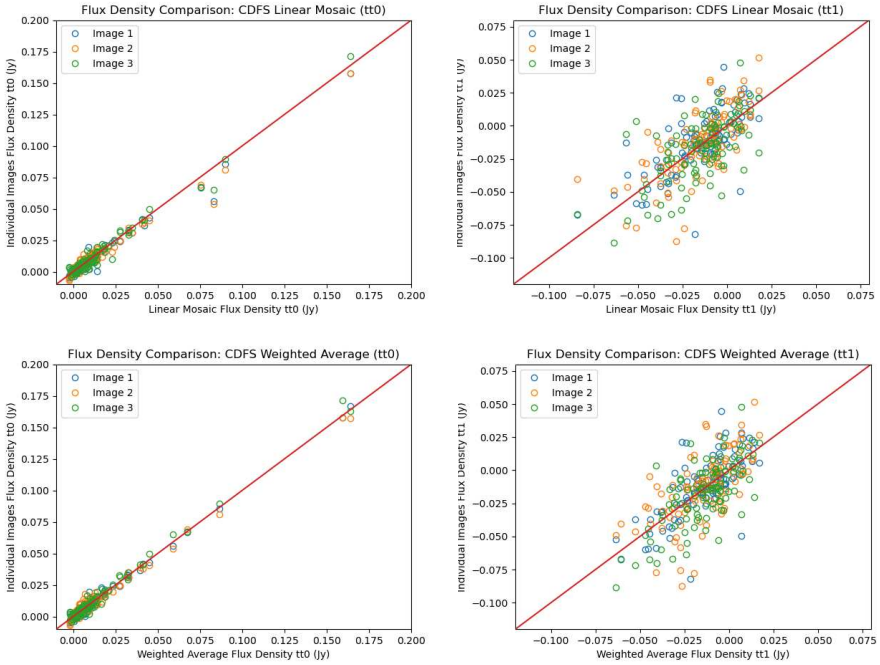

Figure 17: Flux density comparisons of the joint imaging combinations results versus the individual input images for CDFS. **Top:** Linear mosaic, **Bottom:** rms weighted average. The left and right columns are the tt0 and tt1 image products respectively. The line is flux equality and not a fit to the data.

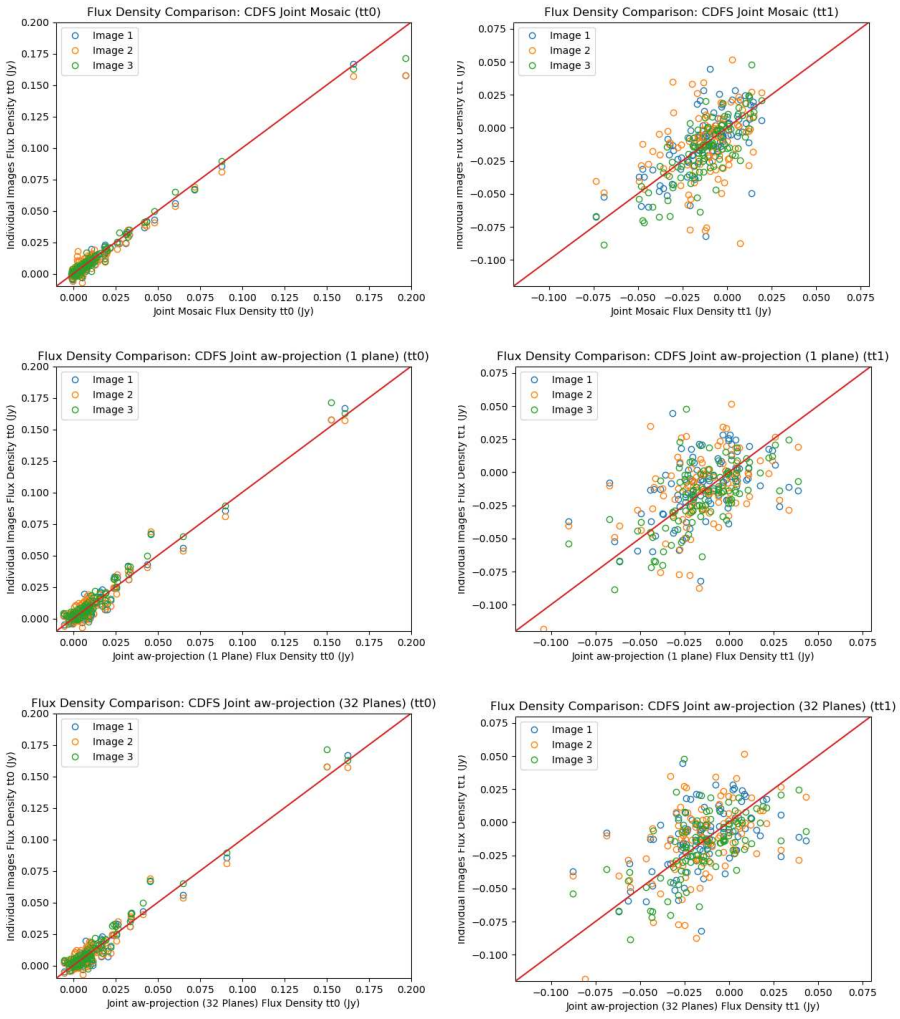

Figure 18: Flux density comparisons of the joint imaging combinations results versus the individual input images for CDFS. **Top:** Mosaic gridder, **Middle:** aw-projection with 1 w-projection plane **Bottom:** aw-projection with 32 w-projection planes. The left and right columns are the tt0 and tt1 image products respectively. The line is flux equality and not a fit to the data.