ALMA Mosaic Imaging Issues Prior to Cycle 6 (CASA 5.4.0)

NAASC Memo # 117

Authors: North American ALMA Science Center (NAASC) Software Support Team & the CASA Team

Date: 28 September, 2018

ABSTRACT

Based on an ALMA helpdesk ticket submitted during Cycle 5 by a very experienced CASA user (i.e. who understood that they were seeing something inconsistent between CASA versions), and subsequent investigations by expert NAASC scientific staff, two issues have been discovered that affect past ALMA mosaic images. It is notable that the mosaic dataset that engendered the ticket and investigation had properties that made these issues particularly apparent (a 15% effect in part of their mosaic). For the majority of ALMA mosaics, the issues are not that large, nor are they obvious by visual inspection, but that is not to say they are not impactful. To assess the severity of the issues we have performed simulations in CASA for 7m-array and 12m-array mosaics with rectangular aspect ratios ranging from 1.1 to 4.2, and analyzed the differences in the resulting images. As a byproduct of this work, we also demonstrate that the flux recovery of simulated sources in mosaic fields is very reliable in CASA 5.4.0 (the Cycle 6 release).

ALMA Mosaic Imaging Issues Prior to Cycle 6 (CASA 5.4.0)

North American ALMA Science Center (NAASC) Software Support Team & the CASA Team

September 28, 2018

Abstract

Based on an ALMA helpdesk ticket submitted during Cycle 5 by a very experienced CASA user (i.e. who understood that they were seeing something inconsistent between CASA versions), and subsequent investigations by expert NAASC scientific staff, two issues have been discovered that affect past ALMA mosaic images. It is notable that the mosaic dataset that engendered the ticket and investigation had properties that made these issues particularly apparent (a 15% effect in part of their mosaic). For the majority of ALMA mosaics, the issues are not that large, nor are they obvious by visual inspection, but that is not to say they are not impactful. To assess the severity of the issues we have performed simulations in CASA for 7m-array and 12m-array mosaics with rectangular aspect ratios ranging from 1.1 to 4.2, and analyzed the differences in the resulting images. As a byproduct of this work, we also demonstrate that the flux recovery of simulated sources in mosaic fields is very reliable in CASA 5.4.0 (the Cycle 6 release).

Mosaic Issue 1:
Scope: Affects all ALMA 7m-array mosaics imaged in CASA prior to Cycle 5 (prior to CASA 5.1.1).
Summary: Prior to Cycle 5, CASA ALMA 7m-array mosaic imaging used a non-optimal definition for the antenna primary beam (PB). For Nyquist-sampled mosaics, after primary beam correction, within the 50% mosaic primary beam response, the result is flux densities that are too high by ∼10% (in =< CASA 4.7.2, i.e. <= Cycle 4), using `clean` and `tclean`. At the 20% mosaic primary beam response level, the flux densities are too large by ∼ 30%.

Mosaic Issue 2:
Scope: Affects all ALMA mosaics (7m and 12m) imaged in CASA during Cycle 5 (CASA 5.1.1).
Summary: In CASA 5.1.1 (Cycle 5), a truncated version of the PB model was erroneously used for gridding the data. This truncation causes aliasing (at some level) in *all* CASA mosaic images, regardless of mosaic shape, or antenna size. More rectangular mosaic patterns have larger discrepancies than more square patterns. However, the manifestation of the problem in the image plane depends sensitively on a range of factors including: mosaic shape, mosaic spatial sampling, and uniformity of mosaic sensitivity per pointing. In *extreme* cases errors as large as ±15% have been observed inside the 50% mosaic primary beam response (i.e. the PI data that started our investigation). However, the median error inside the 50% mosaic response for the simulated data presented here is less than 2.5%, even for an aspect ratio of 4. That said, the situation rapidly deteriorates for more rectangular patterns, such that between the 0.2 and 0.5 PB levels, flux density errors as large as 30% can occur for aspect ratios $\gtrsim 1.5$.

The simulations developed for this study will be used in future Cycles to ensure these issues (or new ones) cannot go unnoticed. In summary, we note that though much consternation has been generated about the mosaic issues in the Cycle 5 version of CASA 5.1.1-5, inside the 0.5 primary beam response pattern, the Cycles 1-4 issues with 7m-mosaics is a significantly larger effect in general. Furthermore, the absence of accurate ALMA beam models will continue to result in significant flux density errors in the outer portions of ALMA images, of order 12% at the 0.2 PB level.

Contents

1 Introduction

Based on an ALMA helpdesk ticket submitted during Cycle 5 by a very experienced CASA user, and subsequent investigations, two issues have been discovered that affect past ALMA mosaic images. It is notable that the mosaic that engendered the ticket and investigation had properties that made these issues particularly apparent (15% in part of their mosaic). For the majority of ALMA mosaics, the issues are not that large, nor very obvious by visual inspection, but that is not to say, not impactful. This memo describes the nature of the two major issues that were discovered, and demonstrates the effects via simulated data. Both issues have been fixed and carefully validated in the CASA 5.4.0 release for ALMA Cycle 6. In §5 we describe other discoveries and changes that were made to improve mosaicing for Cycle 6, and in §6 we describe the deficiencies of the current ALMA beam models in CASA.

For reference, Table 1 summarizes the CASA versions used in ALMA production as a function of ALMA Cycle. For the tests described in this memo, CASA 5.4.0-68 was used to test the status of CASA for ALMA Cycle 6 imaging, CASA 5.1.1-5 was used to test the Cycle 5 behavior, and CASA 4.7.2 (Cycle 4) was used to test the behavior of CASA for ALMA Cycles 1-4. In all these cases, the Red Hat Enterprise Linux 6 version of CASA was used on the NAASC cluster. The specific behavior reported here for 7m-array mosaics in CASA 4.7.2 (Cycle 4) was also spot checked back to CASA 4.0 (ALMA Cycle 1).

Table 1: CASA versions used in ALMA Production and their corresponding issues

CASA version	ALMA Cycle	Mosaic issue 1	Mosaic issue 2
4.0, 4.1[a]	1	Yes	No
4.2.2	2	Yes	No
4.3.1	3	Yes	No
4.4.0[a]	3	Yes	No
4.5.1[b]	3	Yes	No
4.5.2[c]	3	Yes	No
4.7.0	4	Yes	No
4.7.2[b]	4	Yes	No
5.1.1-5	5	No	Yes
5.4.0-68	6	No	No

[a]only used for manual processing
[b]mid-Cycle planned pipeline patch
[c]pipeline bugfix release

2 Simulations, Imaging, and Analysis

We have chosen simulations as the primary means of demonstrating these issues, but we have also verified that the trends reported here are reproduced for representative real datasets from the Cycle 6 ALMA pipeline benchmark ensemble of 94 Member ObsUnitSets (MOUS; that list was given to DMG as part of the C6 pipeline delivery). Simulated images are preferable for characterization because it is possible to put a dense grid of sources throughout an image both at set locations and more randomly, as well as to make unambiguous comparisons of the model flux density versus the achieved flux density. The 7m-array and 12m-array simulations were carried out in the CASA 5.4.0 software, and each traverse the mosaic 4 times with 30 second integrations for a total of 2 minutes on-source time per mosaic pointing. (For the small mosaics this would be a single execution, for the largest ones it would correspond to several execution blocks). The simulations used a Nyquist-sampled hexagonal grid pattern of pointings (ranging from 9 to 144 pointings in a variety of aspect ratios, see Table 2). Only realistic thermal noise is present (from the receiver and atmosphere), no residual phase noise was added. Configuration 43-3 was used to simulate the 12m-array data, and the standard ACA configuration was used to simulate the 7m-array data. For the 12m and 7m-array simulations, the sky model was constructed from 0.02 and 0.2 mJy beam^{-1} point sources, respectively. The bandwidth of the simulations is 1 GHz and the center frequency is 100 GHz. The simulations are centered at transit for declination $-53°$, with the longest being almost 5 hr and spanning a symmetric elevation range of 49 to 60°. Since point source visibilities are calculated by direct Fourier transform in CASA, there are no potential gridding effects introduced at the simulation stage, and the gridding of tclean itself can be cleanly tested.

The simulated mosaic images were made with `gridder='mosaic'`, and `specmode='mfs'`. We chose `imsize` and `cell` size values that are representative of the choices that the ALMA interferometric Imaging pipeline would make in terms of imaging to the 20% mosaic primary beam response level with adequate padding, oversampling the synthesized beam by ∼5 pixels across the minor axis, and employing Briggs `weighting` with `robust=0.5`. Images were first made in CASA 5.4.0 using automasking, and an appropriate rms-based clean threshold. The mask generated from CASA 5.4.0, and the same clean threshold were then used to clean the 5.1.1 and 4.7.2 images. Because use of the pointing table in construction of mosaic images has been turned off by default in CASA 5.4.0 (`usepointing=False`), the mosaic pointing tables were cleared before imaging to mimic this behavior in the earlier versions of CASA (also see § 5.1 of this document).

In order to test CASA 5.4.0 against expectations, Figures 1 and 2 compare the recovered flux density from Gaussian fitting (`imfit`) each simulated source in the mosaics to the expected simulated values, for the 12m and 7m-array data, respectively. These figures demonstrate that the accuracy of flux recovery in CASA 5.4.0 for both arrays is very good.

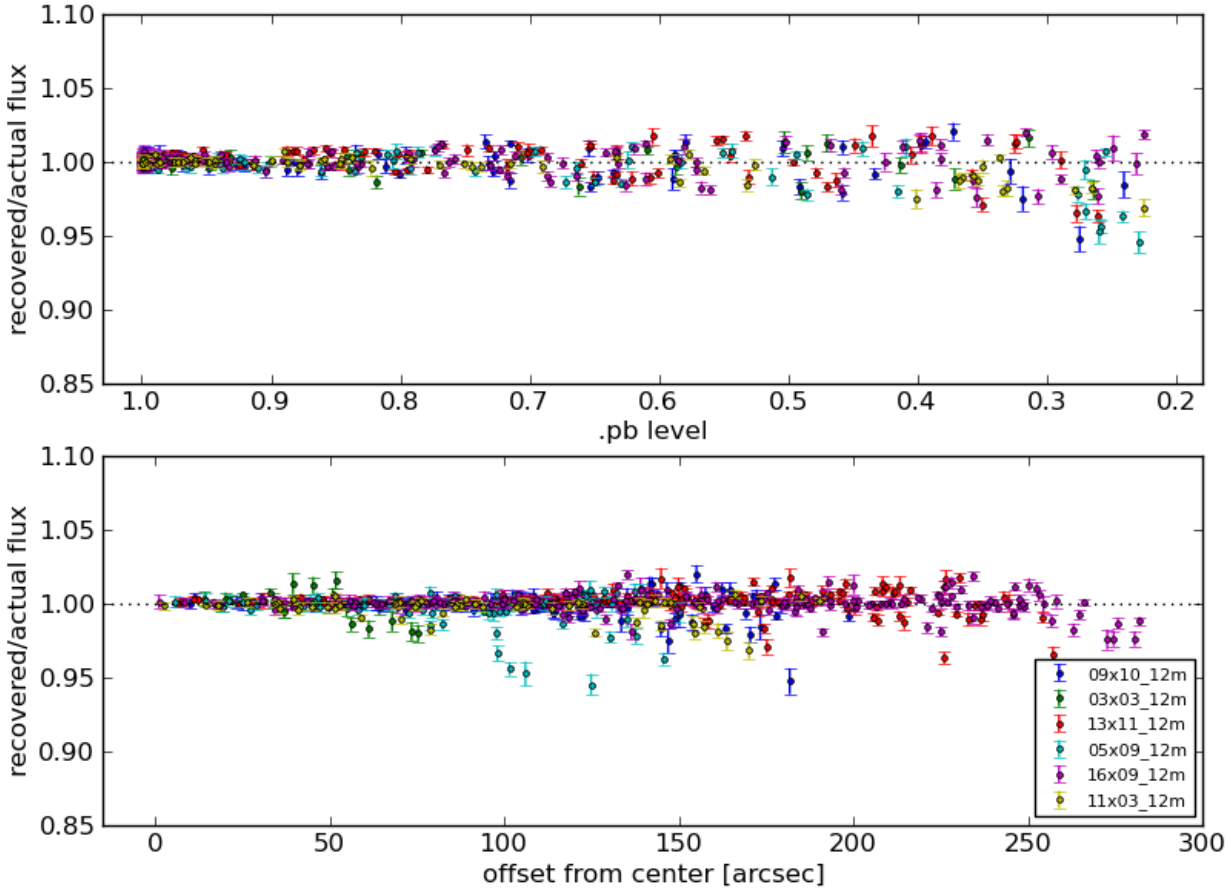

Figure 1: For 12m images created in CASA 5.4.0, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the mosaic center (bottom).

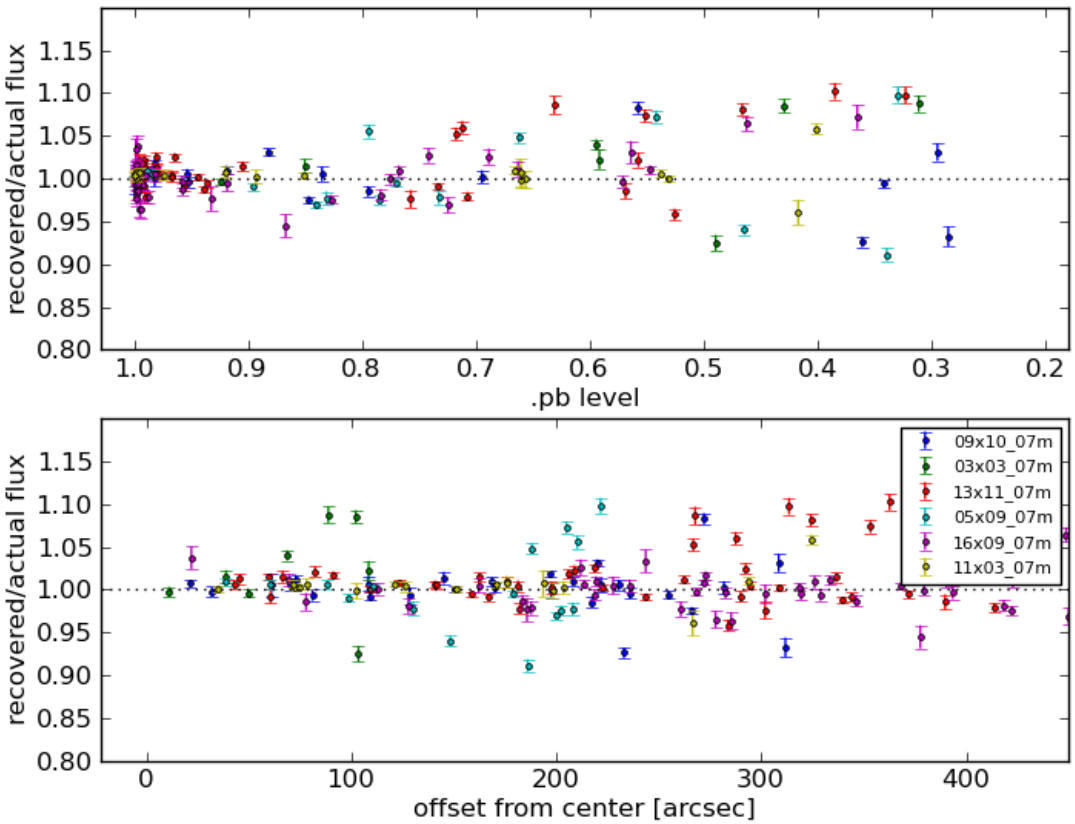

Figure 2: For 7m images created in CASA 5.4.0, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the mosaic center (bottom). The poorer uv coverage of the ACA naturally results in greater scatter in recovered flux densities compared to 12m images.

3 Mosaic Issue 1 (7m array Mosaics Prior to Cycle 5)

This issue affects all ALMA 7m-array mosaics created in CASA up to and including CASA 4.7.2 (Cycles 0 to 4) in both `clean` and when it became available in CASA 4.7.2 (Cycle 4), `tclean`.

3.1 Nature of Mosaic Issue 1

The default primary beam (PB) model for an antenna in CASA is an azimuthally-symmetric (i.e. one-dimensional) Airy disk that would be produced by a telescope with a circular aperture equal to the antenna diameter. The diameter is read from the ANTENNA table of the data, while the central blockage is coded within CASA as a function of the observatory (because this value is not stored in the data). The model is truncated just before the first Airy null (around the 1% response level, i.e. -20 dB). In ALMA Cycle 0, this model was used for all ALMA images, single field or mosaics. However, this default model is not a good representation of the ALMA -10 dB illumination taper which produces a broader beam by $\sim$12% in terms of full width at half power (FWHP) compared to an Airy disk with the antenna's diameter. It is notable that an Airy disk with diameter equal to the antenna diameter is a reasonable approximation for the Very Large Array (VLA) shaped reflectors, which is how it came to be the CASA default.

Before Cycle 1, the disagreement between the assumed CASA model and the design of the ALMA dishes was recognized, and the primary beams in CASA were changed to use a modified diameter Airy disk model with an aperture diameter of 10.7m for the 12m antennas and 6.25m for the 7m antennas. These changes were implemented in the Cycle 1 release version of CASA 4.0. The value of the 0.75 m diameter central blockage was unchanged. These models are in reasonable agreement with the TICRA optical ray tracing models that were commissioned by ALMA during construction to evaluate the antennas and receiver optics. The FWHP of the 12m-antenna Airy disk model matches the average DA/DV antenna TICRA model FWHP to better than 2%. The modified Airy disk diameter for the 7m array antennas, for which TICRA models do not exist, were assumed to scale in a similar way to the 12m-antennas. Further details on how the agreement between the model beam profiles (TICRA and CASA) and the true antenna beam pattern degrades beyond the FWHP radius are presented in § 6.

Unfortunately, due to a coding error, the modified diameter was not being used specifically for ALMA 7m array mosaics in all CASA versions up to and including CASA 4.7.2 (Cycle 4), instead it would silently fall back to the original CASA default (antenna diameter) for this use case. Based on archaeology performed on old JIRA tickets, when the changes to the modified diameters were made, single field images were carefully tested (and correct for both 7m and 12m). Of course the testers would not have been able to guess that the beam models being used could be different for single fields and mosaics, and as demonstrated in § 3.2, visually the differences are not obvious. This particular issue had been unknowingly fixed by related code changes by CASA 5.1.1 – the Cycle 5 release. The issue is present in both clean and `tclean` prior to CASA 5.1.1.

3.2 Simulated Demonstration of Issue 1

An example comparison of a 7m-array mosaic imaged in CASA 4.7.2 versus CASA 5.4.0 is shown in Fig. 3. Inside the 50% level of the primary beam response, the 4.7.2 image has flux densities that are too high by 10%. This effect can also be seen in the comparison of the percent difference of primary beam response patterns in the bottom right panel of Fig. 3. Statistics for these images are given in Tables 2 and 3. Additional examples of 7m-array simulated mosaic results from CASA 4.7.2 are shown in Appendix A1.

The direct comparison of the simulated point source flux density of each point source to the flux density in the `tclean` mosaic image, Gaussian fitted using `imfit` is shown in Figures 4 and 5 for the 12 and 7m arrays, respectively. These Figures should be directly compared to Figures 1 and 2 for images created in CASA 5.4.0. The dominant effect is the systematic offset of flux density for all sources in 7m images. No serious issues are apparent in the 12m-array mosaics in 4.7.2, though there is some evidence for minor degradation in the recovered flux in the 0.2 to 0.5 PB range, greater scatter is also apparent compared to 5.4.0.

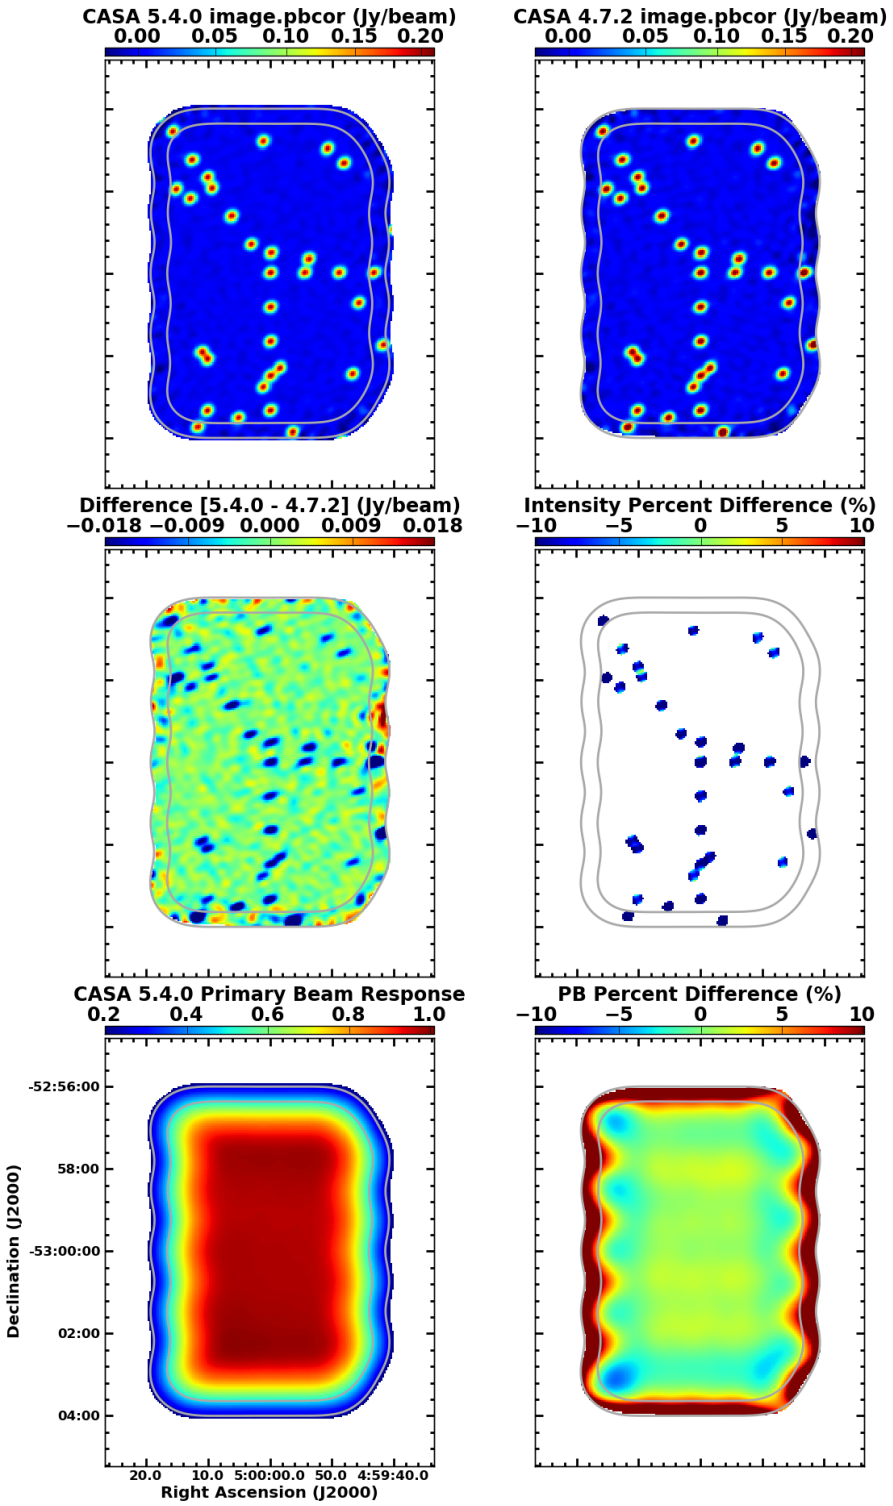

Figure 3: Demonstration of Mosaic issue 1 in CASA 4.7.2. **Top row:** `tclean` images of a 7m array simulation with 5×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam response (`.pb`) image from 5.4.0 and the percentage difference from the corresponding `.pb` image from 4.7.2. On all panels, the grey contours show the 0.25 and 0.5 levels of the primary beam response. Statistics of these images are given in Tables 2 and 3. The $\sim 10\%$ too high flux densities (inside the 0.5 .pb level) that arise from CASA 4.7.2 for 7m-array data is readily apparent.

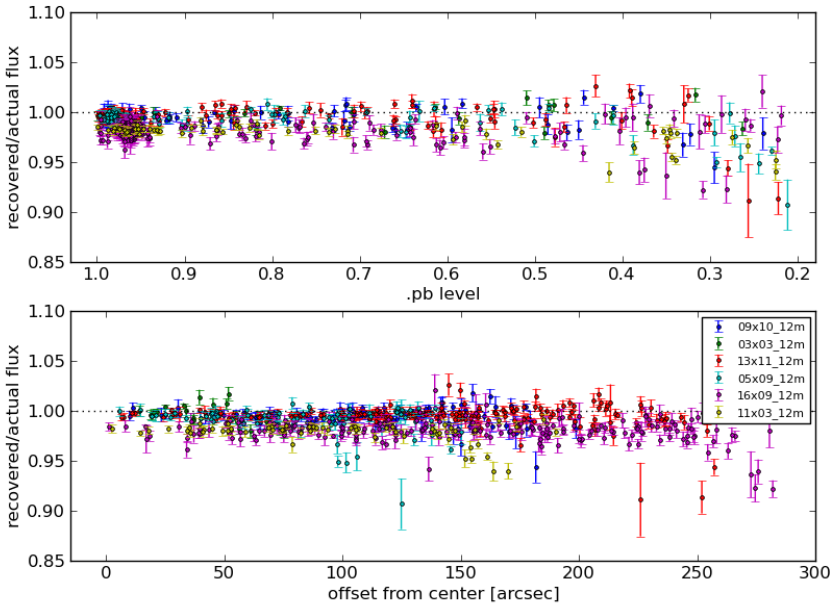

Figure 4: For 12m images created in CASA 4.7.2, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the center of the mosaic (bottom). The identical ratio for images in CASA 5.4.0 is plotted in Figure 1.

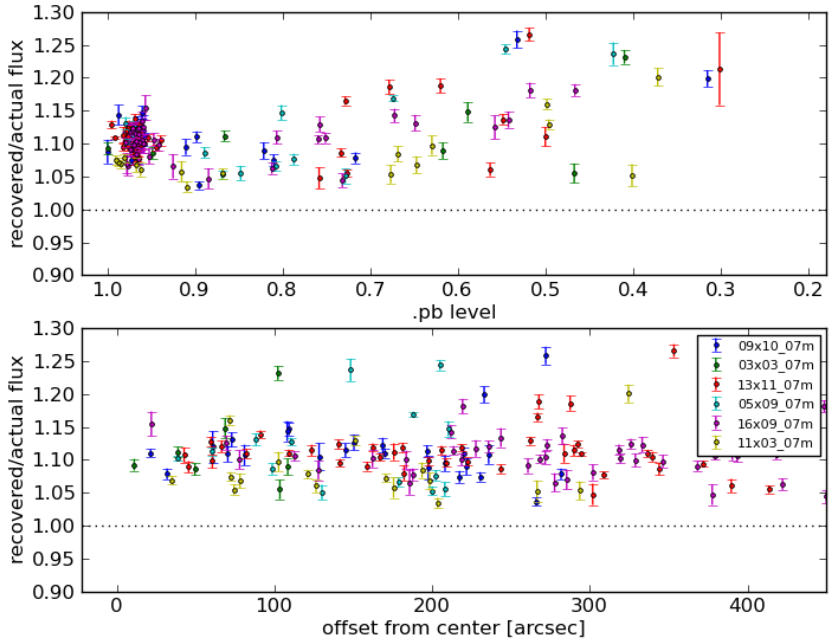

Figure 5: For 7m images created in CASA 4.7.2, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the center of the mosaic (bottom). The poorer uv coverage of the ACA naturally results in greater scatter in recovered flux densities compared to 12m images. The identical ratio for images in CASA 5.4.0 is plotted in Figure 2.

4 Mosaic Issue 2 (All Mosaics Cycle 5)

Affects all ALMA mosaics imaged in CASA `tclean` in the Cycle 5 pipeline.

4.1 Nature of Mosaic Issue 2

CASA 5.1.1 (actually CASA 5.0, but this was not a production version of CASA for ALMA) underwent a major infrastructure change known as the switch from VI1 to VB2. As part of this major change many parts of the CASA imaging code had to be modified/re-implemented. Due to an error introduced during the overhaul, a truncated version of the primary beam model (the modified Airy disk models described in § 3.1) was used for gridding the data in CASA 5.1.1. This truncation causes aliasing (at some level) in all CASA mosaic images, regardless of mosaic shape, or antenna size. However, the effects become larger the more elongated (rectangular) the mosaic pattern is. Note it does not depend on whether the `imsize` is square, but how square the sky coverage pattern is. It is also notable that no mosaic patterns produced automatically by the ALMA OT are 100% square, there are only relatively square and progressively less square patterns, including completely non-square patterns. Elongation of the mosaic pattern north-south versus east-west produce similar levels of error. For more rectangular patterns, the aliasing visibly manifests as artifacts or "wings" outside the nominal 20% mosaic primary beam response (see Fig. 6).

Indeed, such aliasing artifacts outside the nominal 20% mosaic primary beam response (PB) were reported by the pipeline working group for highly non-square mosaic images prior to release of the Cycle 5 pipeline but it was thought to only affect that particular use case, and to only affect the region outside the nominal 0.2 PB level. Unfortunately because this one of a large number of issues that came about due to the infrastructure change, and at the time did not seem as serious as some others, it was not prioritized for developer effort.

It was only during the course of investigation of the aforementioned helpdesk ticket late in Cycle 5 that the nature of the issue was fully understood and thus its broader impact – i.e. affecting the imaged flux densities inside the 20% and even the 50% primary beam response levels for extreme cases. The manifestation of the problem depends sensitively on a range of factors including: mosaic shape, mosaic spatial sampling, and uniformity of mosaic sensitivity per pointing. In extreme cases, errors as large as ±5% inside the 0.5 mosaic primary beam response have been observed, but the median values are less than 2.5% even for aspect ratios of ~ 4. Between the 0.2 to 0.5 PB levels, the errors can be as large as 30% for highly rectangular patterns.

4.2 Simulated Demonstration of Issue 2

An example comparison of a 7m-array mosaic imaged in CASA 5.1.1 versus CASA 5.4.0 is shown in Fig. 6. This Nyquist sampled mosaic has a 5×9 (axis ratio 1.8) pointing pattern. The Intensity percent difference image demonstrates that within the 50% primary beam response, the flux densities are within a few percent of the correct values, with a median of +6.4%, while in the 20% to 50% region the difference is as large as +16.3% (see Table 3). We also made a 7m-array simulation with 5×9 pointings but using the NRAO VLA Sky Survey (NVSS) pattern, i.e. beams overlapping by only FWHM/$\sqrt{2}$, and the resulting differences were similar in magnitude to the Nyquist sampled case (see Table 3).

The direct comparison of the simulated point source flux density of each point source to the flux density in the `tclean` mosaic image, Gaussian fitted using `imfit` is shown in Figures 7 and 8 for the 12 and 7m arrays, respectively. These Figures should be directly compared to Figures 1 and 2 for images created in CASA 5.4.0. The largest effect is the systematic offset for the most elongated mosaics, in the 0.2 to 0.5 PB regime.

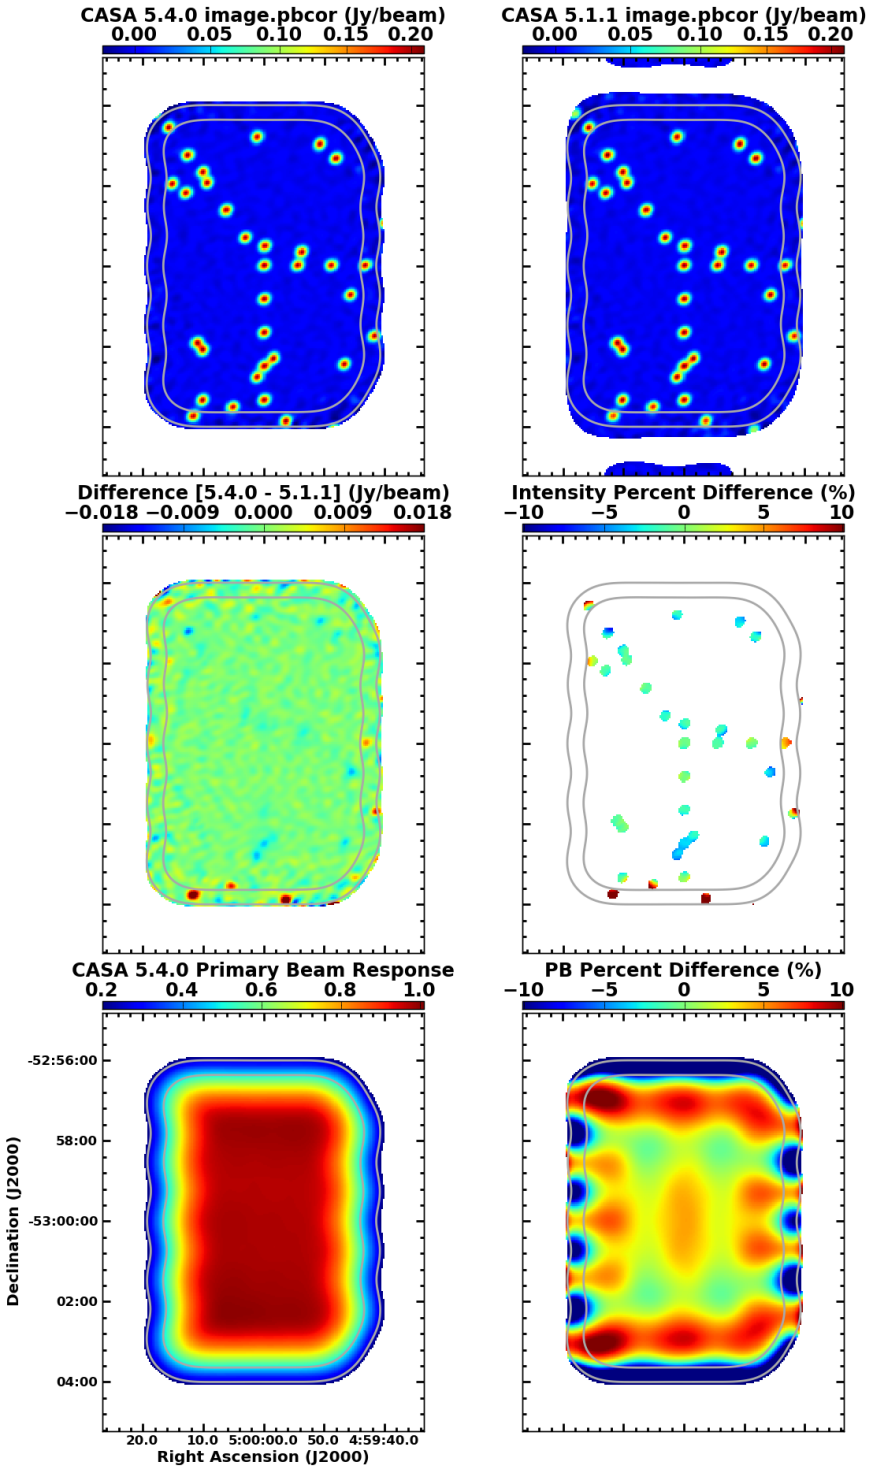

Figure 6: Demonstration of Mosaic issue 2 in CASA 5.1.1. **Top row:** `tclean` images of a 7m array simulation with 5×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from 5.4.0 and the percentage difference from the corresponding `.pb` image from 5.1.1. Statistics of these images are given in Tables 2 and 3. Artifacts caused by the aliasing are visible at the top and bottom.

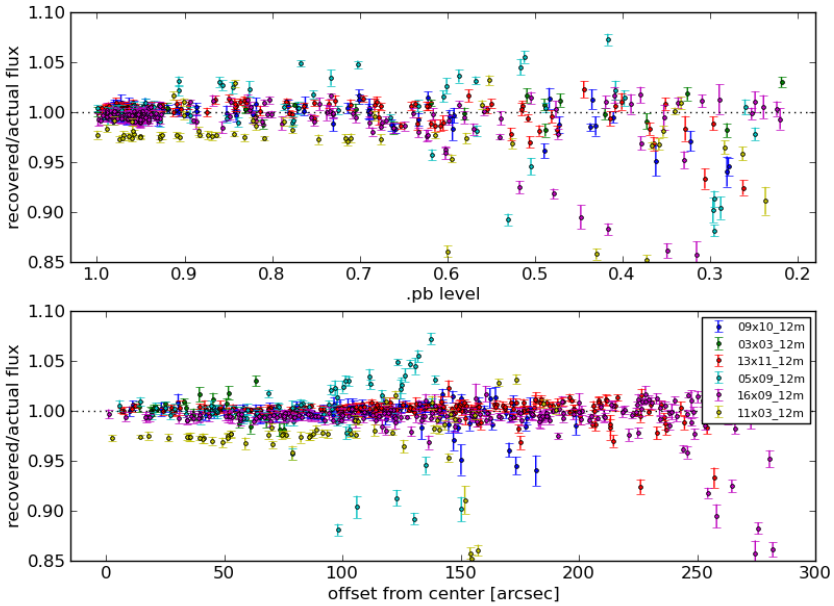

Figure 7: For 12m images created in CASA 5.1.1, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the center of the mosaic (bottom). The identical ratio for images in CASA 5.4.0 is plotted in Figure 1.

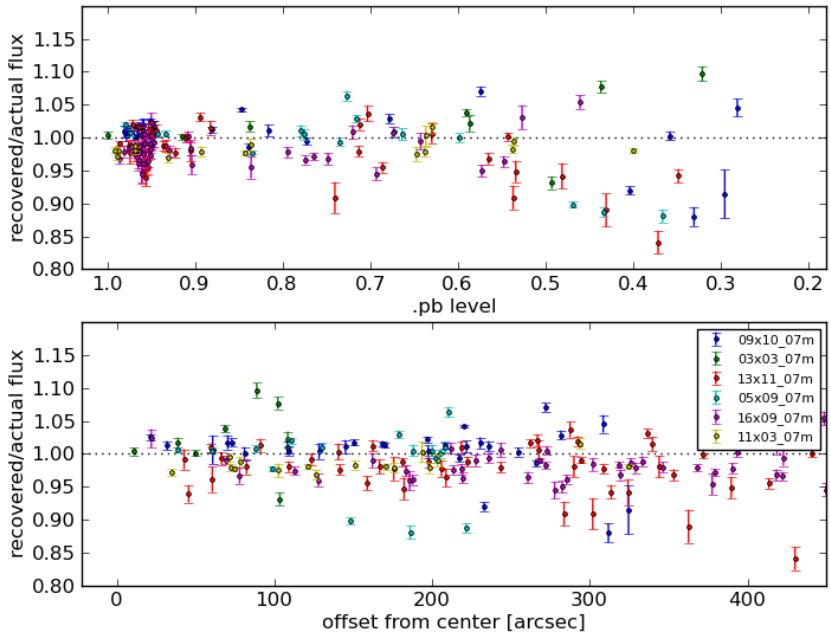

Figure 8: For 7m images created in CASA 5.1.1, fitted point source flux densities divided by the simulated flux densities. This ratio is plotted as a function of primary beam level (top), and distance from the center of the mosaic (bottom). The poorer uv coverage of the ACA naturally results in greater scatter in recovered flux densities compared to 12m images. The identical ratio for images in CASA 5.4.0 is plotted in Figure 2.

Table 2: Statistics of cleaned images and primary-beam-corrected cleaned images from simulations

Antenna diameter	Simulation pattern	Aspect ratio[a]	image rms (mJy beam^{-1}) 4.7.2	5.1.1	5.4.0	pbcor image rms (mJy beam^{-1}) 4.7.2	5.1.1	5.4.0
7m	3x3	1.2	4.36	3.44	3.51	11.53	9.07	9.28
7m	5x9	1.6	2.49	1.76	1.85	6.41	4.60	4.79
7m	9x10	1.1	2.03	1.66	1.78	5.20	4.25	4.56
7m	11x3	4.2	2.40	1.96	1.77	6.25	5.18	4.55
7m	13x11	1.4	1.99	1.42	1.65	5.24	3.67	4.31
7m	16x9	2.1	1.98	1.54	1.75	5.19	3.95	4.54
7m	5x9 NVSS[b]	1.5	2.46	1.86	2.10	6.34	4.73	5.43
12m	3x3	1.2	0.11	0.12	0.12	0.30	0.30	0.30
12m	5x9	1.6	0.12	0.10	0.12	0.31	0.26	0.30
12m	9x10	1.1	0.12	0.11	0.12	0.31	0.29	0.31
12m	11x3	4.2	0.12	0.10	0.12	0.31	0.26	0.30
12m	13x11	1.4	0.12	0.11	0.12	0.32	0.27	0.31
12m	16x9	2.1	0.14	0.11	0.12	0.35	0.27	0.32

[a]Image aspect ratio (this is not simply the ratio of number of pointings since the hexagonal pattern has smaller spacing between rows than along rows).

[b]The spacing of this pattern is coarser: FWHM/$\sqrt{2} = 0.7989\lambda$/D; the rest are 0.51×FWHM = $\lambda/(D\sqrt{3})$.

5 Other issues discovered during this investigation

5.1 Use of the dynamic pointing corrections

While ALMA visibilities are automatically flagged online if either antenna of a given baseline is not within the required tolerance of the intended source direction, small deviations from the desired direction will still be present in the data. In principle, the effect of these deviations can be corrected by using the encoder values from both antennas comprising each baseline. This technique may become important in the highest frequency bands where the primary beam is small, and in particular for single-pointings and mosaic images of areas that contain multiple strong sources. The `tclean` task is able to apply temporal pointing corrections to the visibilities due to pointing errors of the individual telescopes as indicated in the POINTING table of the measurement set. In CASA versions prior to 5.4.0, this capability was enabled by default; in fact, it has been enabled ever since it was introduced. Although ALMA ASDMs do contain values in the Pointing.bin table, which gets filled to the MeasurementSet's POINTING table, they have never been commissioned on an end-to-end basis and the origin of the values are not documented by the observatory. After examining the values for a few datasets, it was immediately clear that they can be incorrect, and we should not be applying them, so we added a new parameter to `tclean` called `usepointing` with a default value of `False`. The VLA pipeline also does not use this feature of `tclean` (the table is cleared if present). As a result, this capability of `tclean` is also not fully commissioned. A future effort will be needed to properly commission the pointing tables for ALMA and VLA observations, as well as their use in `tclean` itself.

The CASA simulator creates a pointing table in RA/Dec, whereas in 5.1.1 specifically, `tclean` unfortunately assumed the table is in Az/El (consistent with real data). Thus the conversion from pointing table frame to imaging frame is wrong, causing serious image degradation for simulated mosaics. Thus, simulated mosaic measurement sets imaged in CASA 5.1.1 should have their POINTING table cleared before imaging. Even though the issue has been fixed in CASA 5.4.0, the pointing table was cleared for all the simulated results reported here to avoid any uncertainty in the tests.

6 Known deviations between the ALMA beam model in CASA and the actual ALMA primary beam

There are many reasons why the current CASA ALMA modified diameter Airy disk primary beam model is not a good representation of the actual effective antenna beam pattern that is imposed upon real astronomical data. This topic has been discussed at length in the single dish imaging memo (Brogan & Hunter 2014, "ALMA Single Dish

Table 3: Statistics of the percentage difference images of the primary-beam-corrected images from pairs of simulations

Antenna diameter	Simulation pattern	Issue#1: 5.4.0 vs 4.7.2						Issue#2: 5.4.0 vs. 5.1.1					
		5%ile (∼min)		95%ile (∼max)		50%ile (median)		5%ile (∼min)		95%ile (∼max)		50%ile (median)	
		inner[a]	outer[b]	inner	outer	inner	outer	inner	outer	inner	outer	inner	outer
7m	3x3	-12.5	-37.1	-8.1	-12.0	-10.2	-17.5	-0.8	-4.3	+0.0	+0.6	-0.5	-0.3
7m	5x9	-17.2	-47.1	-3.1	-11.0	-9.6	-28.8	-4.4	-1.8	+2.0	+16.3	-1.5	+6.4
7m	9x10	-14.2	-47.9	-6.3	-13.6	-10.4	-31.8	-1.8	-4.4	+2.0	+4.7	-0.1	+0.4
7m	11x3	-14.8	-23.5	+3.0	+4.4	-6.1	-13.0	-1.8	+0.4	+5.1	+24.6	+2.5	+15.5
7m	13x11	-13.9	-47.7	-5.9	-13.7	-10.5	-26.8	-1.6	-3.6	+0.9	+9.8	-0.4	+1.6
7m	16x9	-16.2	-35.9	-1.3	-3.9	-8.9	-21.4	-3.6	-1.5	+4.9	+29.8	+0.6	+10.3
7m	5x9 NVSS	-19.8	-49.9	-6.8	-20.0	-12.9	-31.7	-2.7	-7.2	+2.1	+12.9	-0.3	+2.1
12m	3x3	-2.9	-0.1	-2.9	-6.5	-0.2	-0.1	-0.1	-0.3	+0.2	+0.7	+0.0	+0.4
12m	5x9	-2.8	-5.8	+5.0	+7.5	+0.7	+0.6	-2.6	-4.0	+0.9	+14.7	-0.5	+2.6
12m	9x10	-1.9	-7.4	+2.3	+7.2	+0.3	-0.1	-0.6	-1.3	+0.5	+1.9	+0.0	+0.3
12m	11x3	-17.0	-19.8	+11.9	+13.0	+1.0	+0.5	-2.0	-2.9	+4.4	+24.1	+2.4	+3.8
12m	13x11	-4.0	-10.3	+3.3	+7.1	+0.3	+0.0	-1.2	-1.6	+0.8	+5.8	-0.1	+1.2
12m	16x9	-5.3	-13.6	+6.5	+14.4	+1.8	+1.1	-1.2	-2.1	+3.1	+21.7	+1.0	+2.3

[a]inner region defined as the area where the primary beam response level = 0.5 to 1.0

[b]outer annulus defined as the area where the primary beam response level = 0.2 to 0.5

Imaging Parameters, a Comparison of Observed Beams to Predicted Beams, and a Demonstration of the Method for Measuring Jy/K", see ALMA EDM: SCID-90.04.00.00-0001-A-REP). To summarize the issues here, as previously stated the model is truncated just before the first null, whereas the true instantaneous pattern measured by holography scans shows sidelobes beyond the main lobe (Fig. 9). Second, the effect of the feedlegs makes the pattern non-circularly symmetric, particularly in the first sidelobe , but also has an effect within the main lobe. The disturbance of circular symmetry differs between the DA and DV antennas due to the difference in where the feed legs originate with respect to the dish surface. Third, residual pointing and tracking errors will serve to smear the ideal beam. Fourth, due to refocusing optics for example at Band 3, the primary beam patterns do not simply scale linearly in width with inverse frequency in all cases as currently assumed; this deviation is particularly notable at the top end of Band 3. Finally, the thermal environment of the antenna can cause distortions of the beamshape that vary with time. All of these phenomena (except the fourth one) will tend to push the effective beam pattern of an ALMA dataset toward a more Gaussian-shaped main lobe with a slightly wider FWHP.

In order to quantify the expected scale of the these effects on real data, we can compute the difference between the model beam profile and a Gaussian profile as a function of the radius from the beam center. CASA `tclean` produces images out to the 0.2 response level by default, which is also the value of `pblimit` adopted by the ALMA pipeline, so the region above this level denotes the are of primary concern. Fig. 10 shows an overlay of three profiles: (1) the best analytic match to the TICRA model which is a 12m diameter with parabolic illumination pattern with a -12 dB edge taper (using equations in Baars et al. 2003, ALMA Memo 456), (2) the modified Airy profile used in CASA, and (3) a Gaussian profile that matches the FWHP of (1) and (2), which is arguably somewhat optimistic. The two model profiles are in good agreement throughout the main lobe (i.e., above the 0.2 level). The deviation between both models and a Gaussian profile in the inner part of the beam is also small ($< 2\%$). However, the deviation rises in the outer part of the main lobe, where it is $\sim$5% at the 0.3 response level, and $\sim$12% at the edge of the imaged area. Thus, it is important to emphasize that flux density errors of this magnitude (and larger in the higher frequency bands) will continue to persist in the outer areas of ALMA images even though the two software issues discussed earlier in this memo have been rectified as of CASA 5.4.0.

The CASA subsystem scientist has lobbied the ALMA project for the last five years to either commission a complete ensemble of TICRA models for all bands / antenna types or to produce viable primary beam models based on astroholography. To date, requests for the full suite of TICRA models to be commissioned has been denied for cost reasons and, while various astroholography campaigns have been carried out, no observation-based beam models have been communicated to the CASA team.

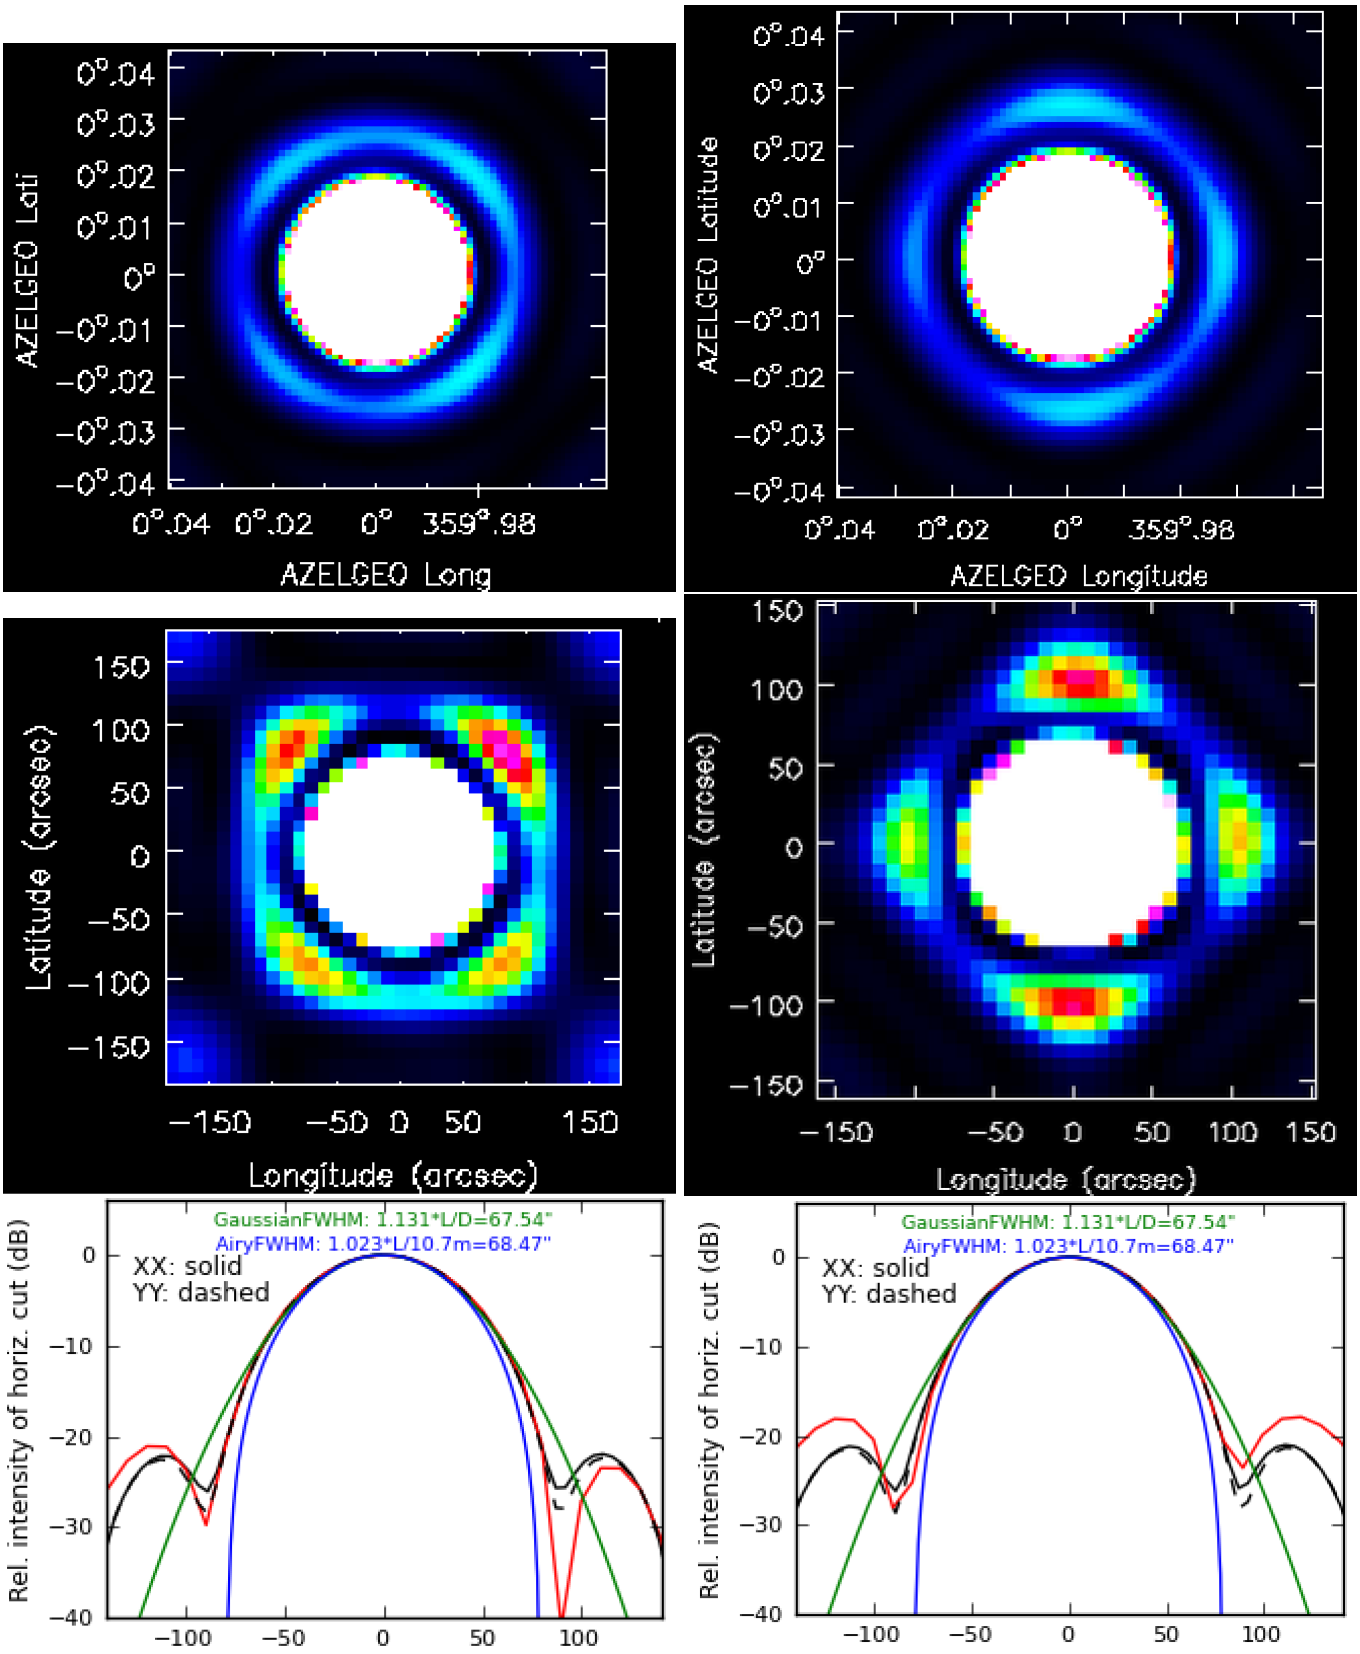

Figure 9: **Left column:** Band 3 beam patterns of DA48: Top image is the TICRA model, middle image is a celestial holography image, bottom shows image cuts with Airy (blue) and Gaussian (green) profiles overlaid. The current CASA model is the blue curve, truncated at -20 dB; **Right column:** Same as left column but for DV17, whose feed legs are oriented at 45° from DA antennas, and which block both the plane wave and the reflected spherical wave.

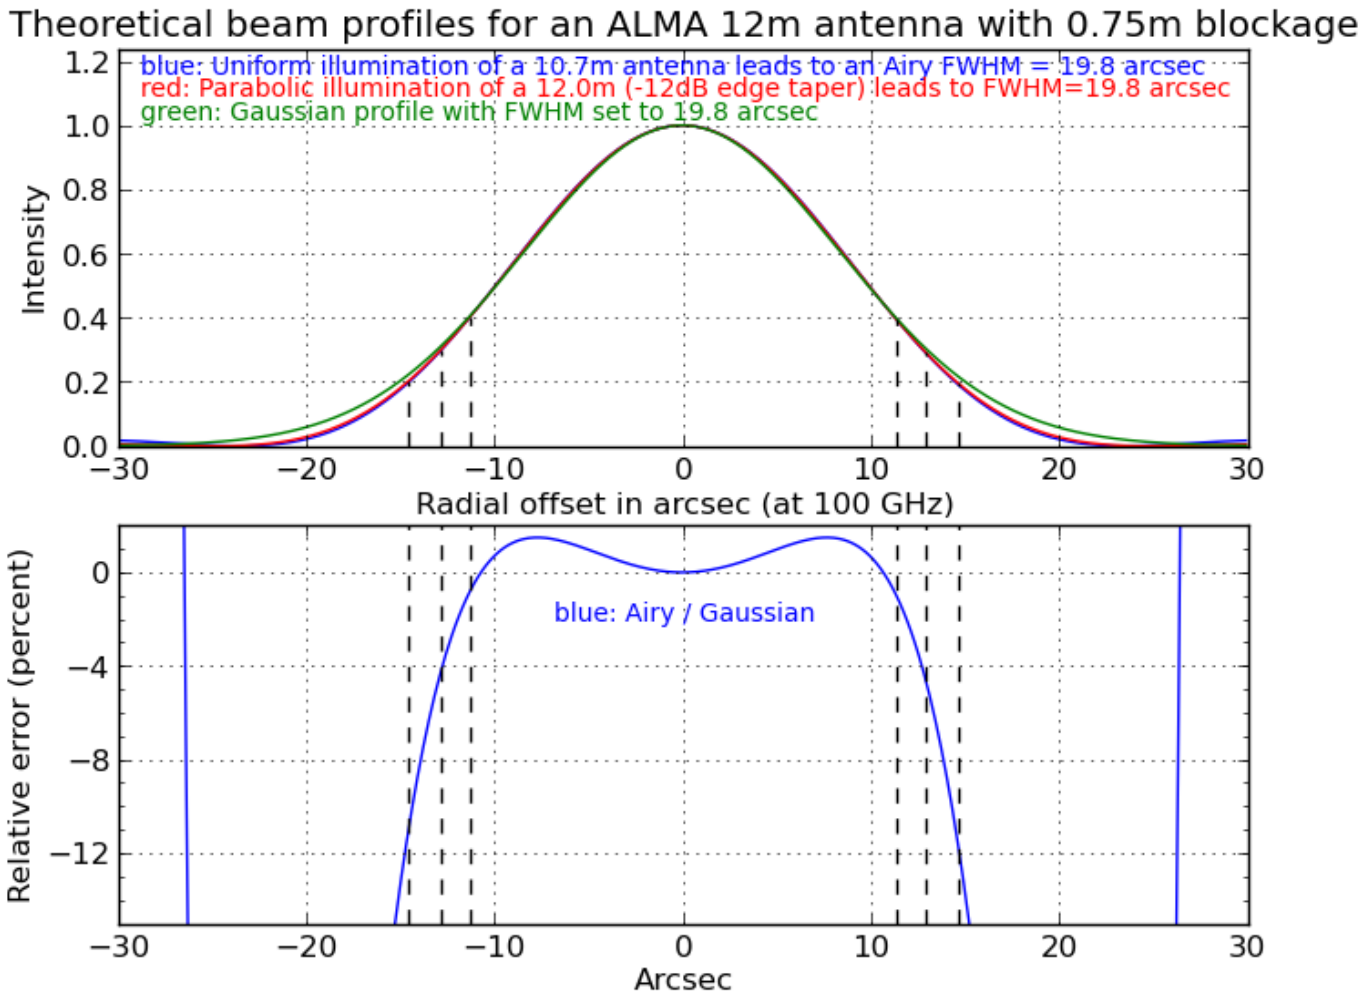

Figure 10: Top panel: analytic model profiles for an ALMA 12m antenna, including parabolic illumination with a -12 dB taper (red curve), a modified Airy profile (blue) for a 10.7m diameter (untruncated), and a Gaussian profile (green) that matches the FWHP diameter of the other two models. Bottom panel: The percentage deviation between the modified Airy model and a Gaussian. Vertical dashed lines denote the 0.4, 0.3 and 0.2 response level in both panels.

Appendices

A Additional examples of mosaic issue 1

A.1 7m 4.7.2 simulations

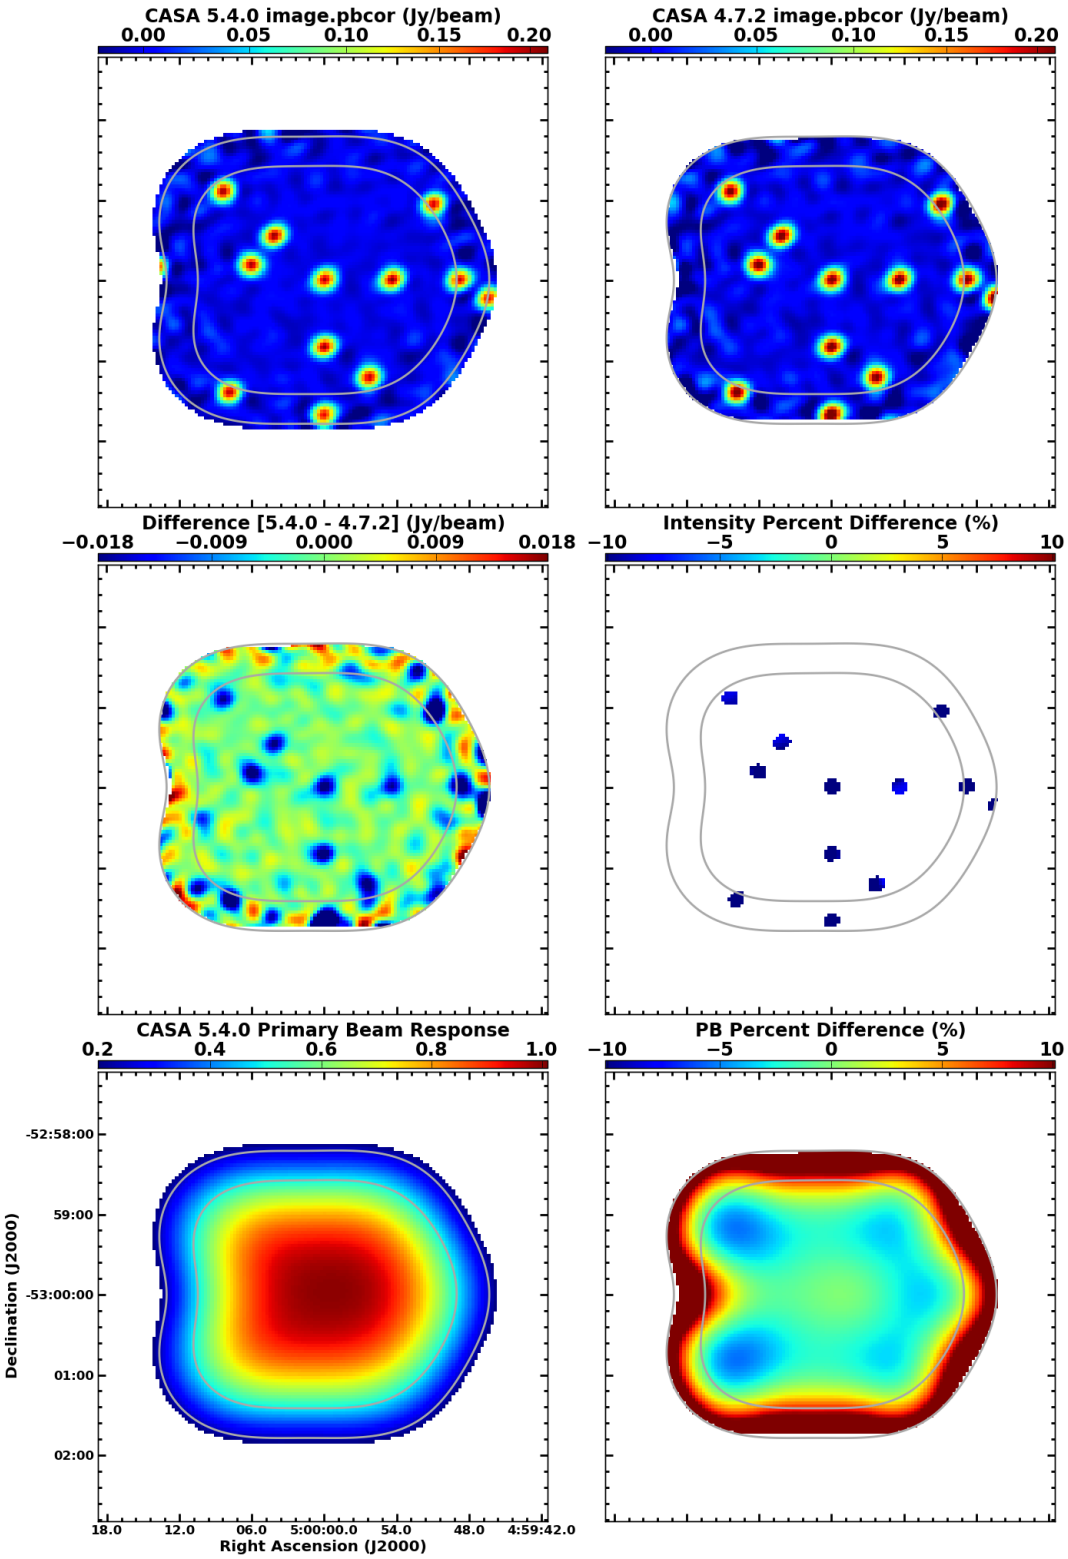

Figure 11: **Top row:** `tclean` images of a 7m array simulation with 3×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 2 of Tables 2 and 3.

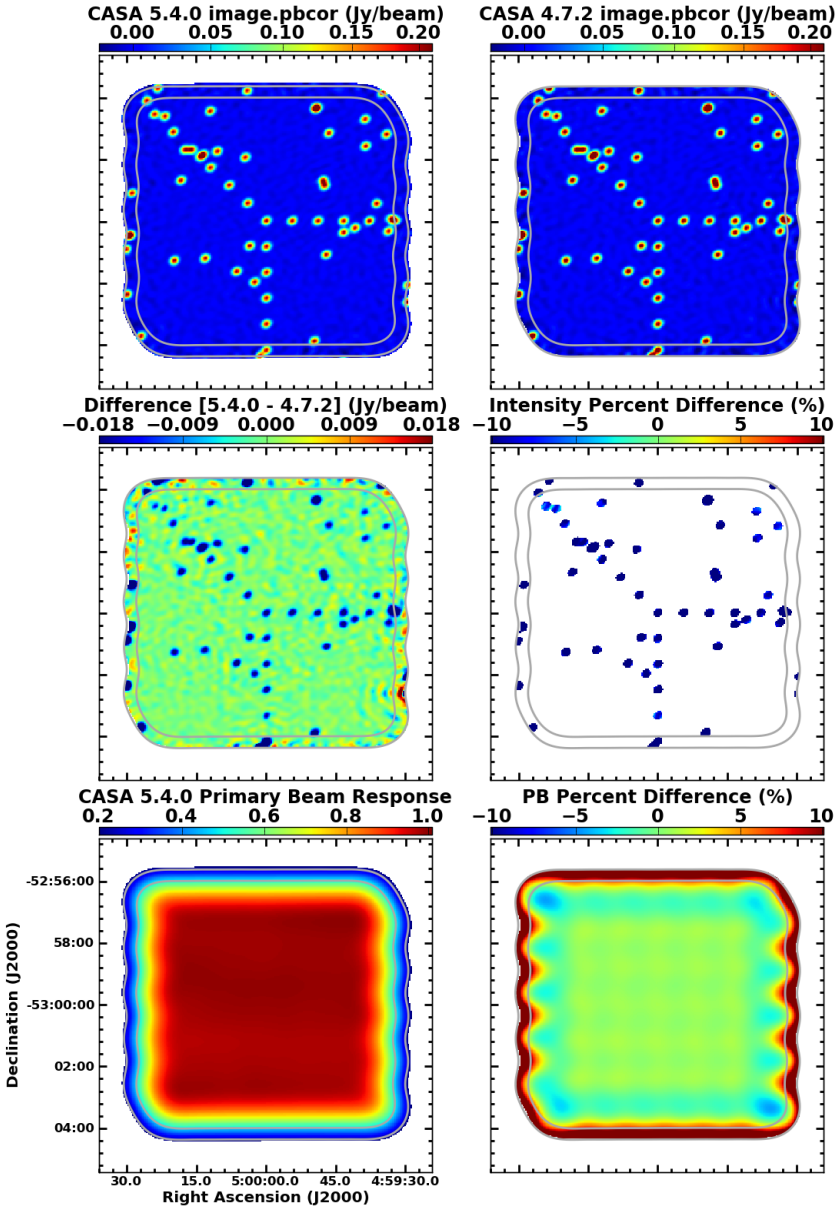

Figure 12: **Top row:** `tclean` images of a 7m array simulation with 9×10 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 3 of Tables 2 and 3.

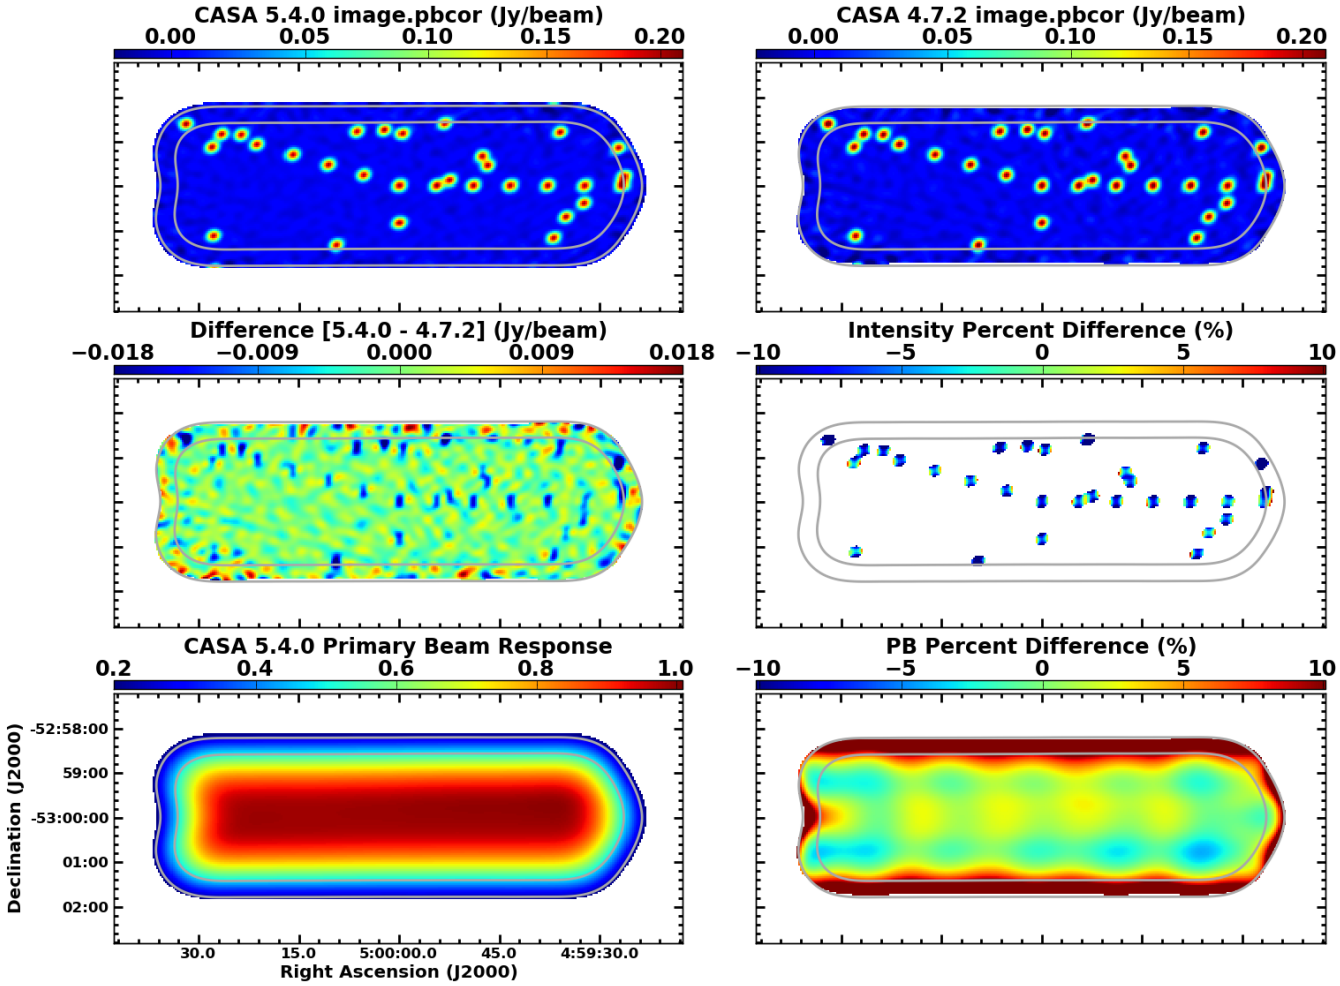

Figure 13: **Top row:** `tclean` images of a 7m array simulation with 11×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 4 of Tables 2 and 3.

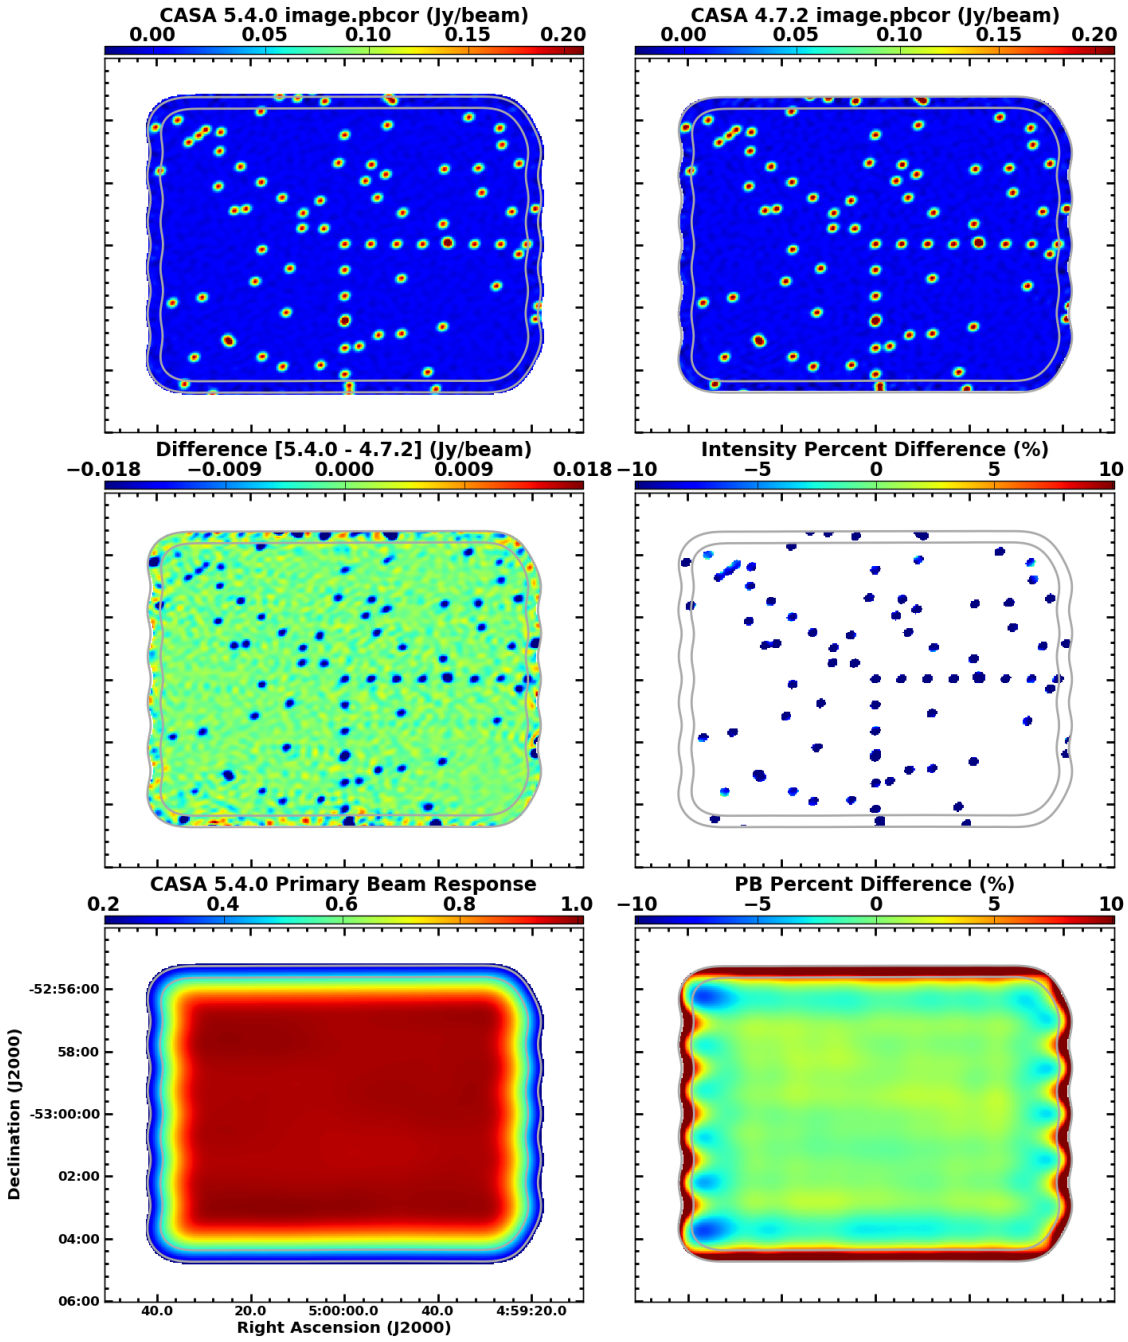

Figure 14: **Top row:** `tclean` images of a 7m array simulation with 13×11 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 5 of Tables 2 and 3.

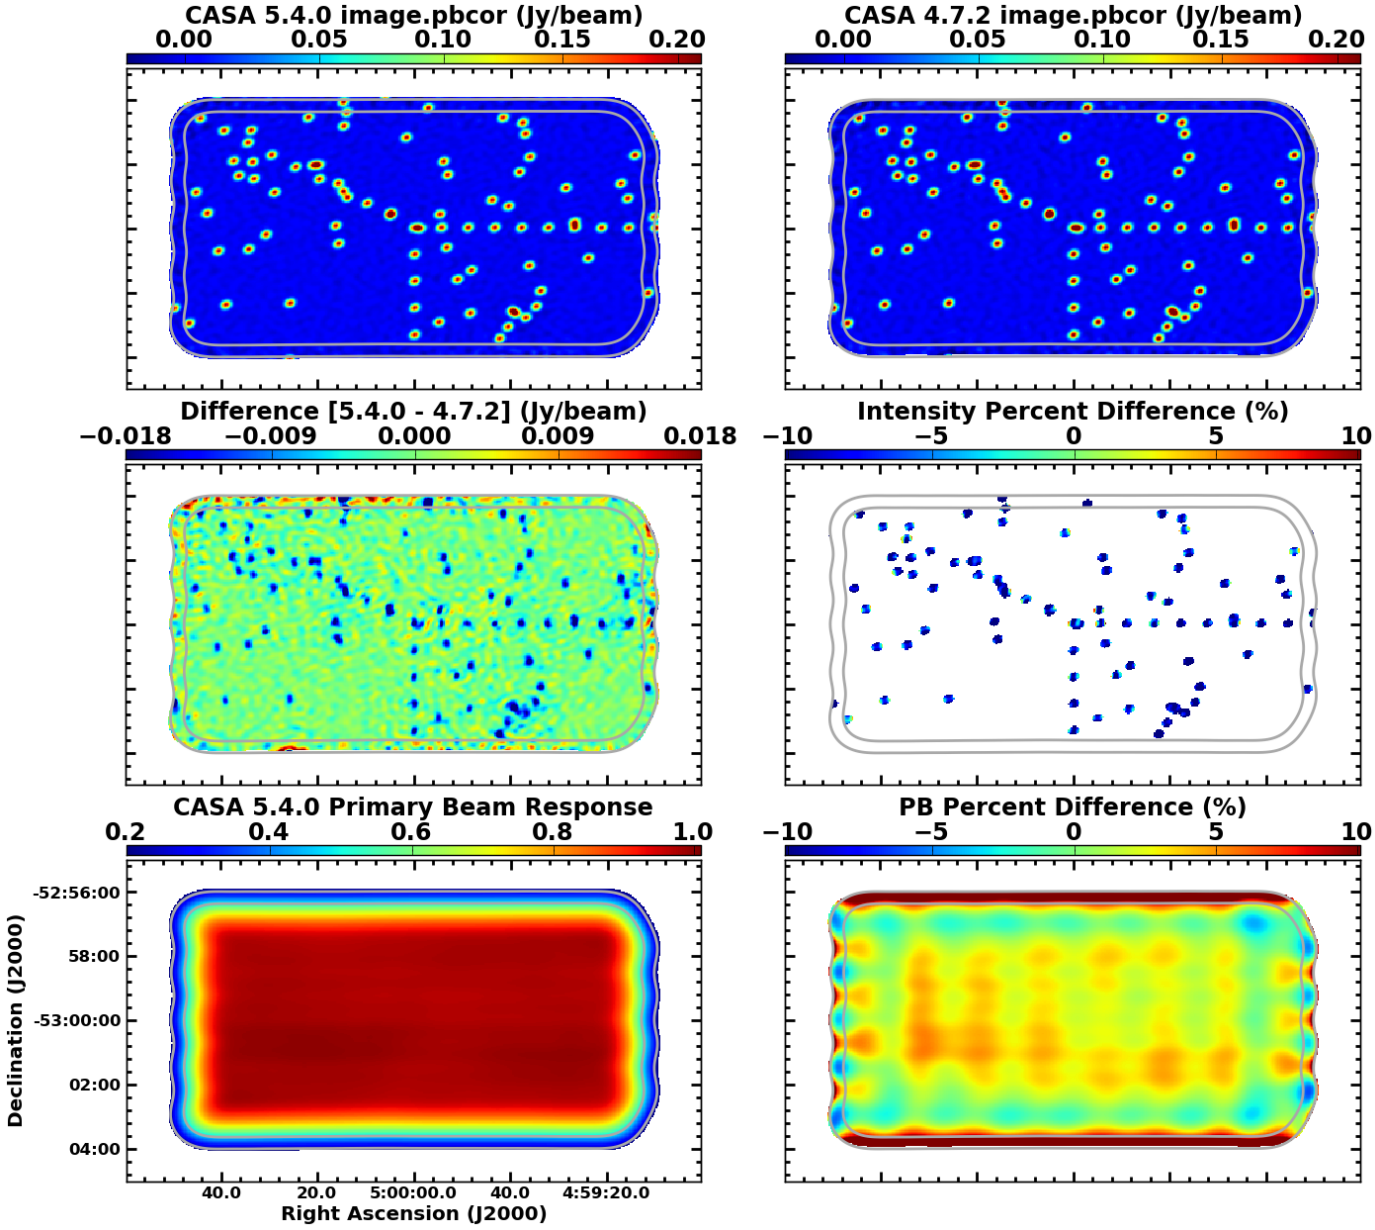

Figure 15: **Top row:** `tclean` images of a 7m array simulation with 16×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 6 of Tables 2 and 3.

A.2 12m 4.7.2 simulations

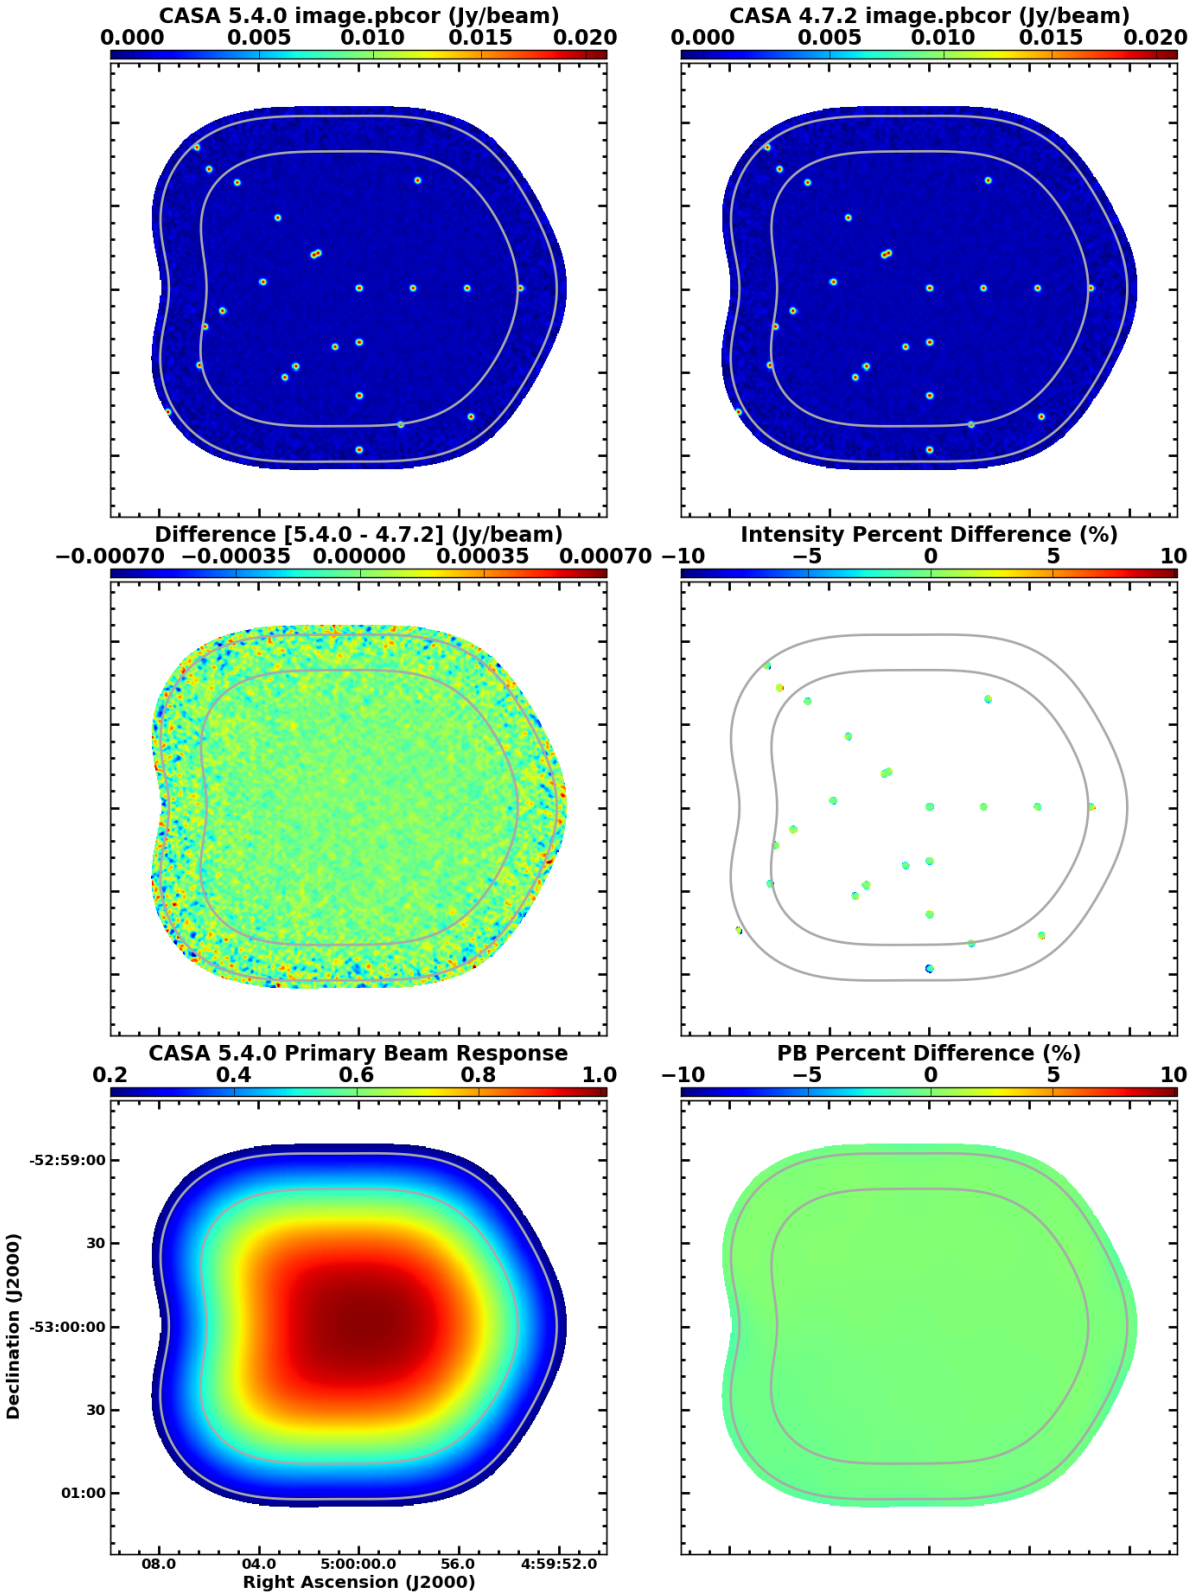

Figure 16: **Top row:** `tclean` images of a 12m array simulation with 3×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 7 of Tables 2 and 3.

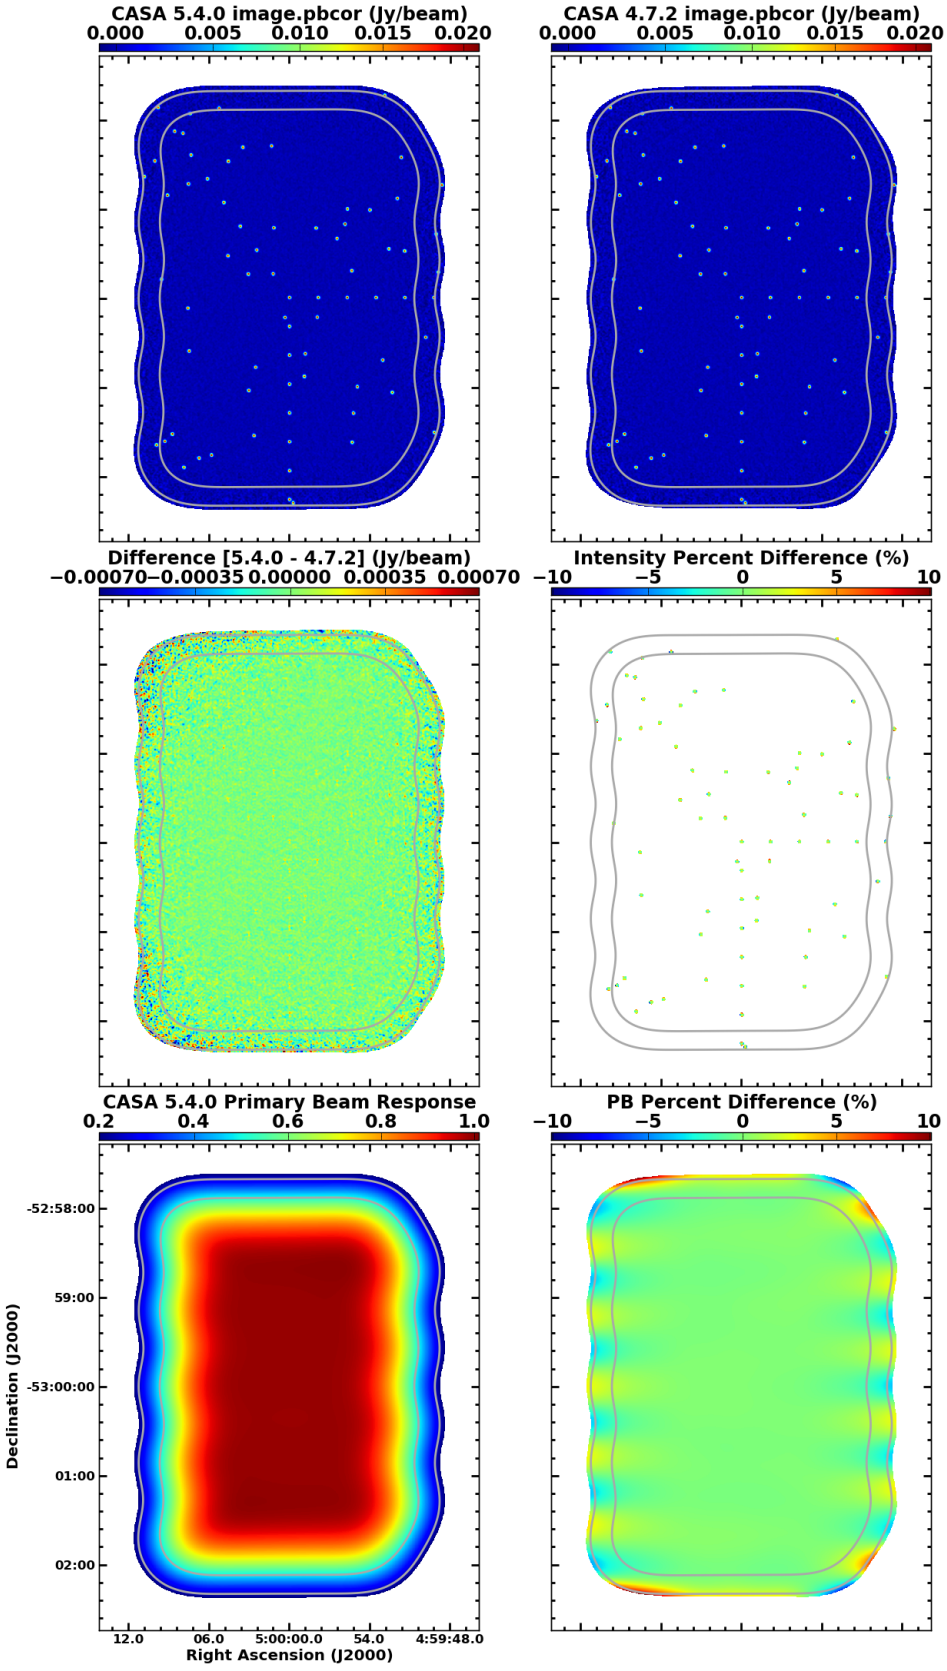

Figure 17: **Top row:** `tclean` images of a 12m array simulation with 5×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 8 of Tables 2 and 3.

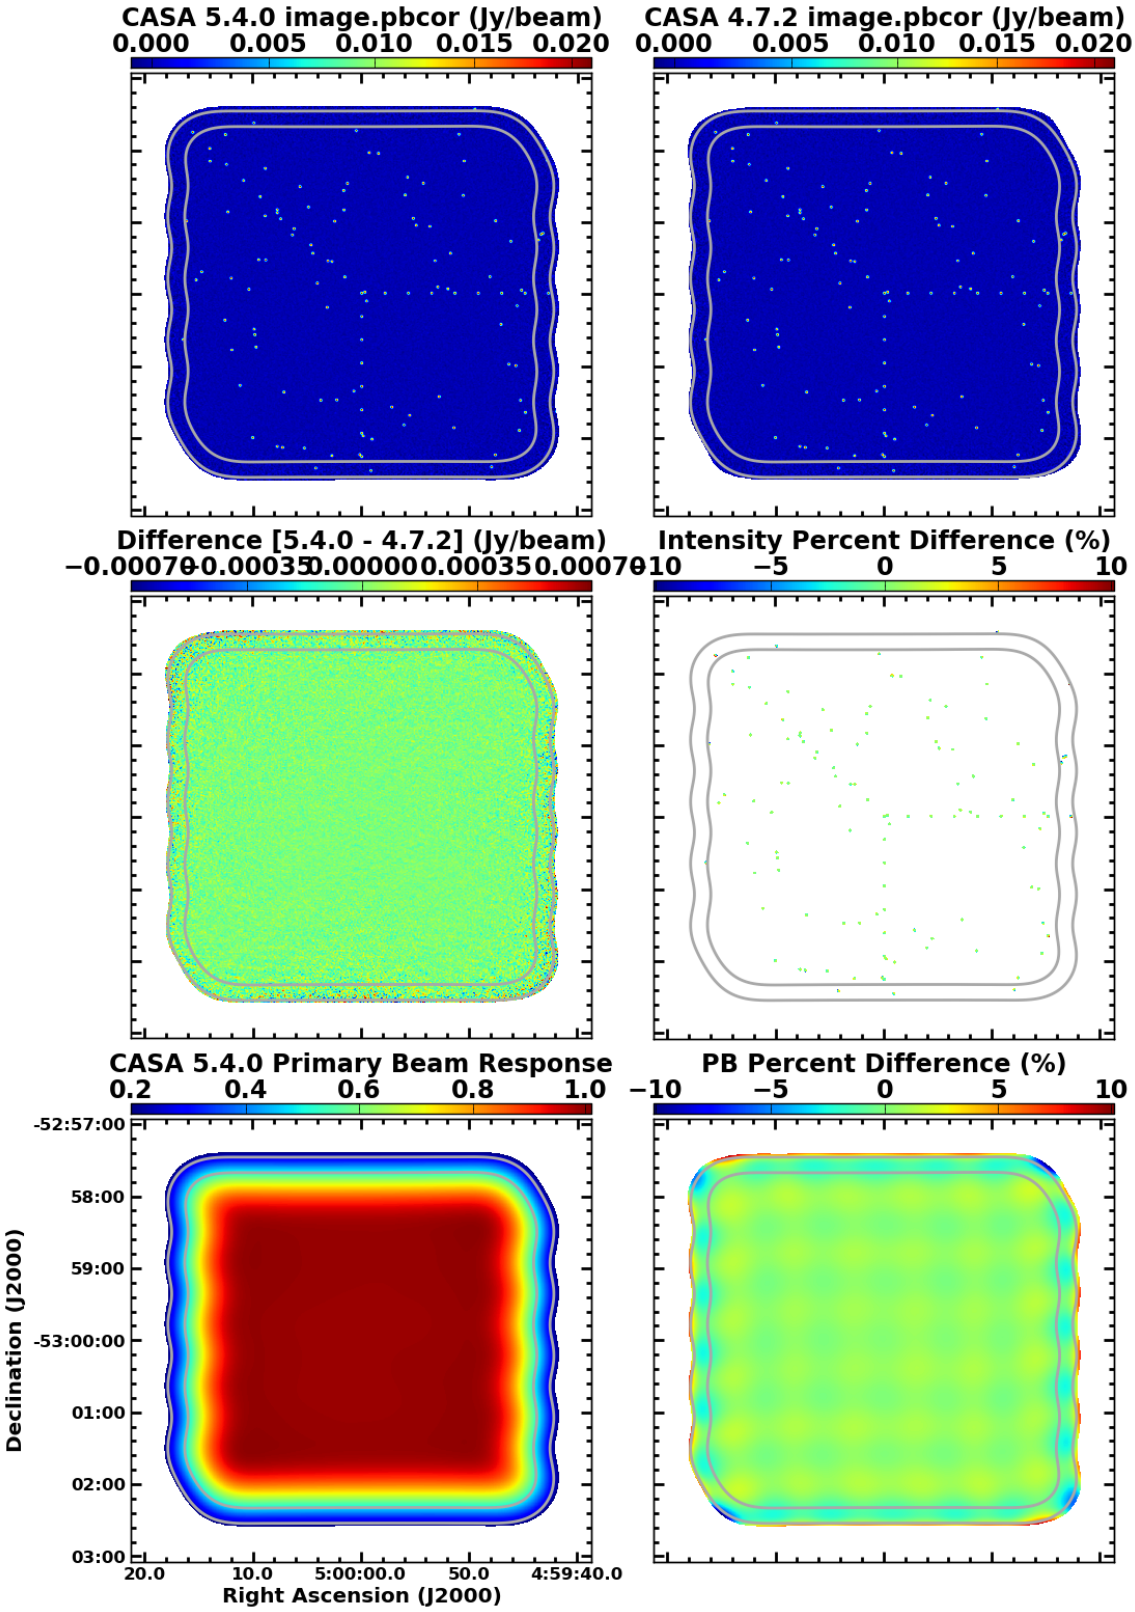

Figure 18: **Top row:** `tclean` images of a 12m array simulation with 9×10 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 9 of Tables 2 and 3.

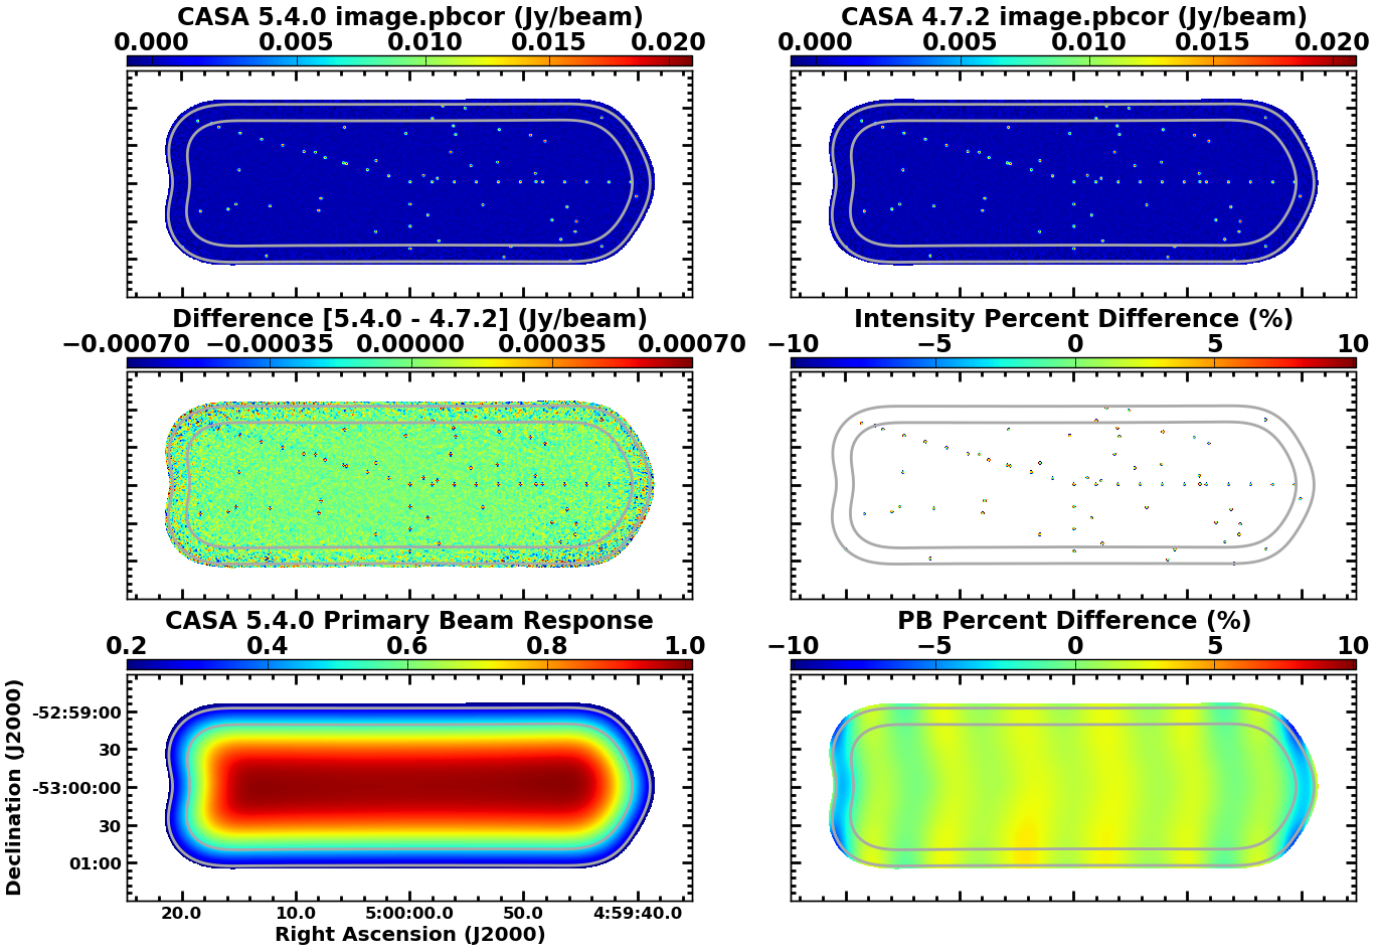

Figure 19: **Top row:** `tclean` images of a 12m array simulation with 11×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 10 of Tables 2 and 3.

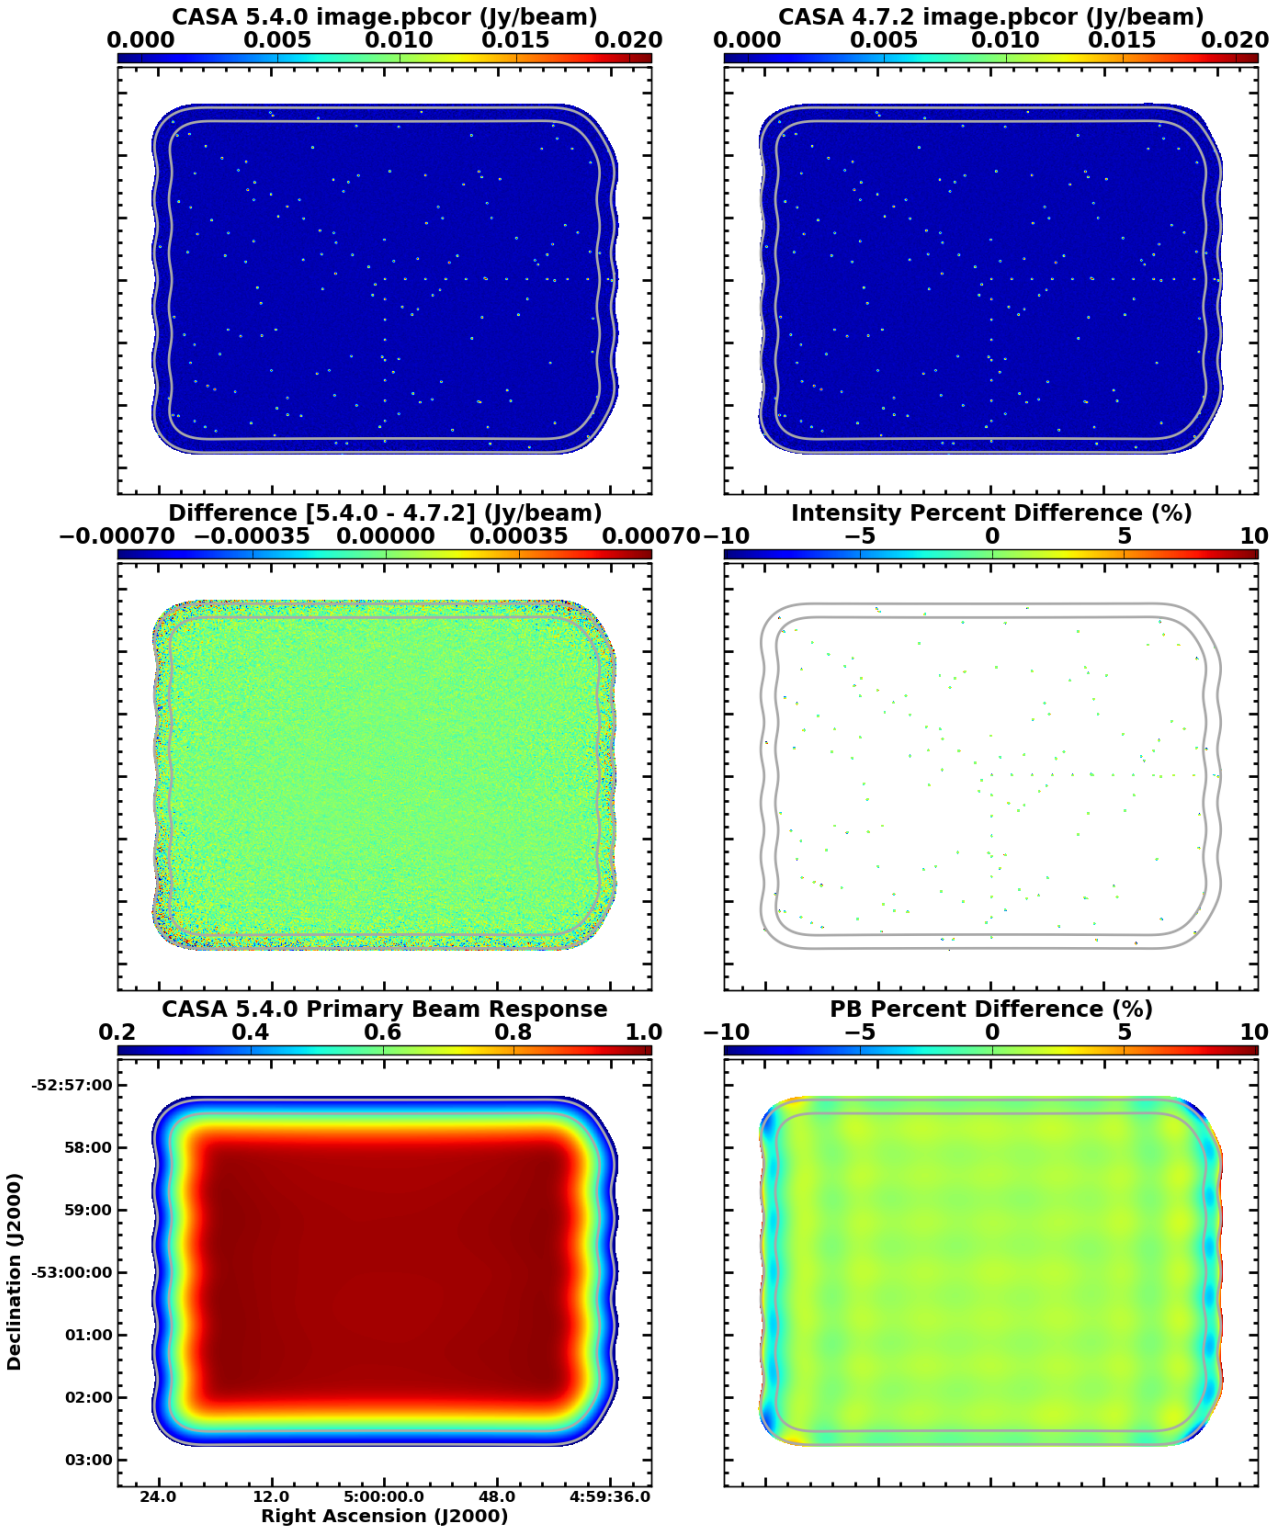

Figure 20: **Top row:** `tclean` images of a 12m array simulation with 13×11 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 11 of Tables 2 and 3.

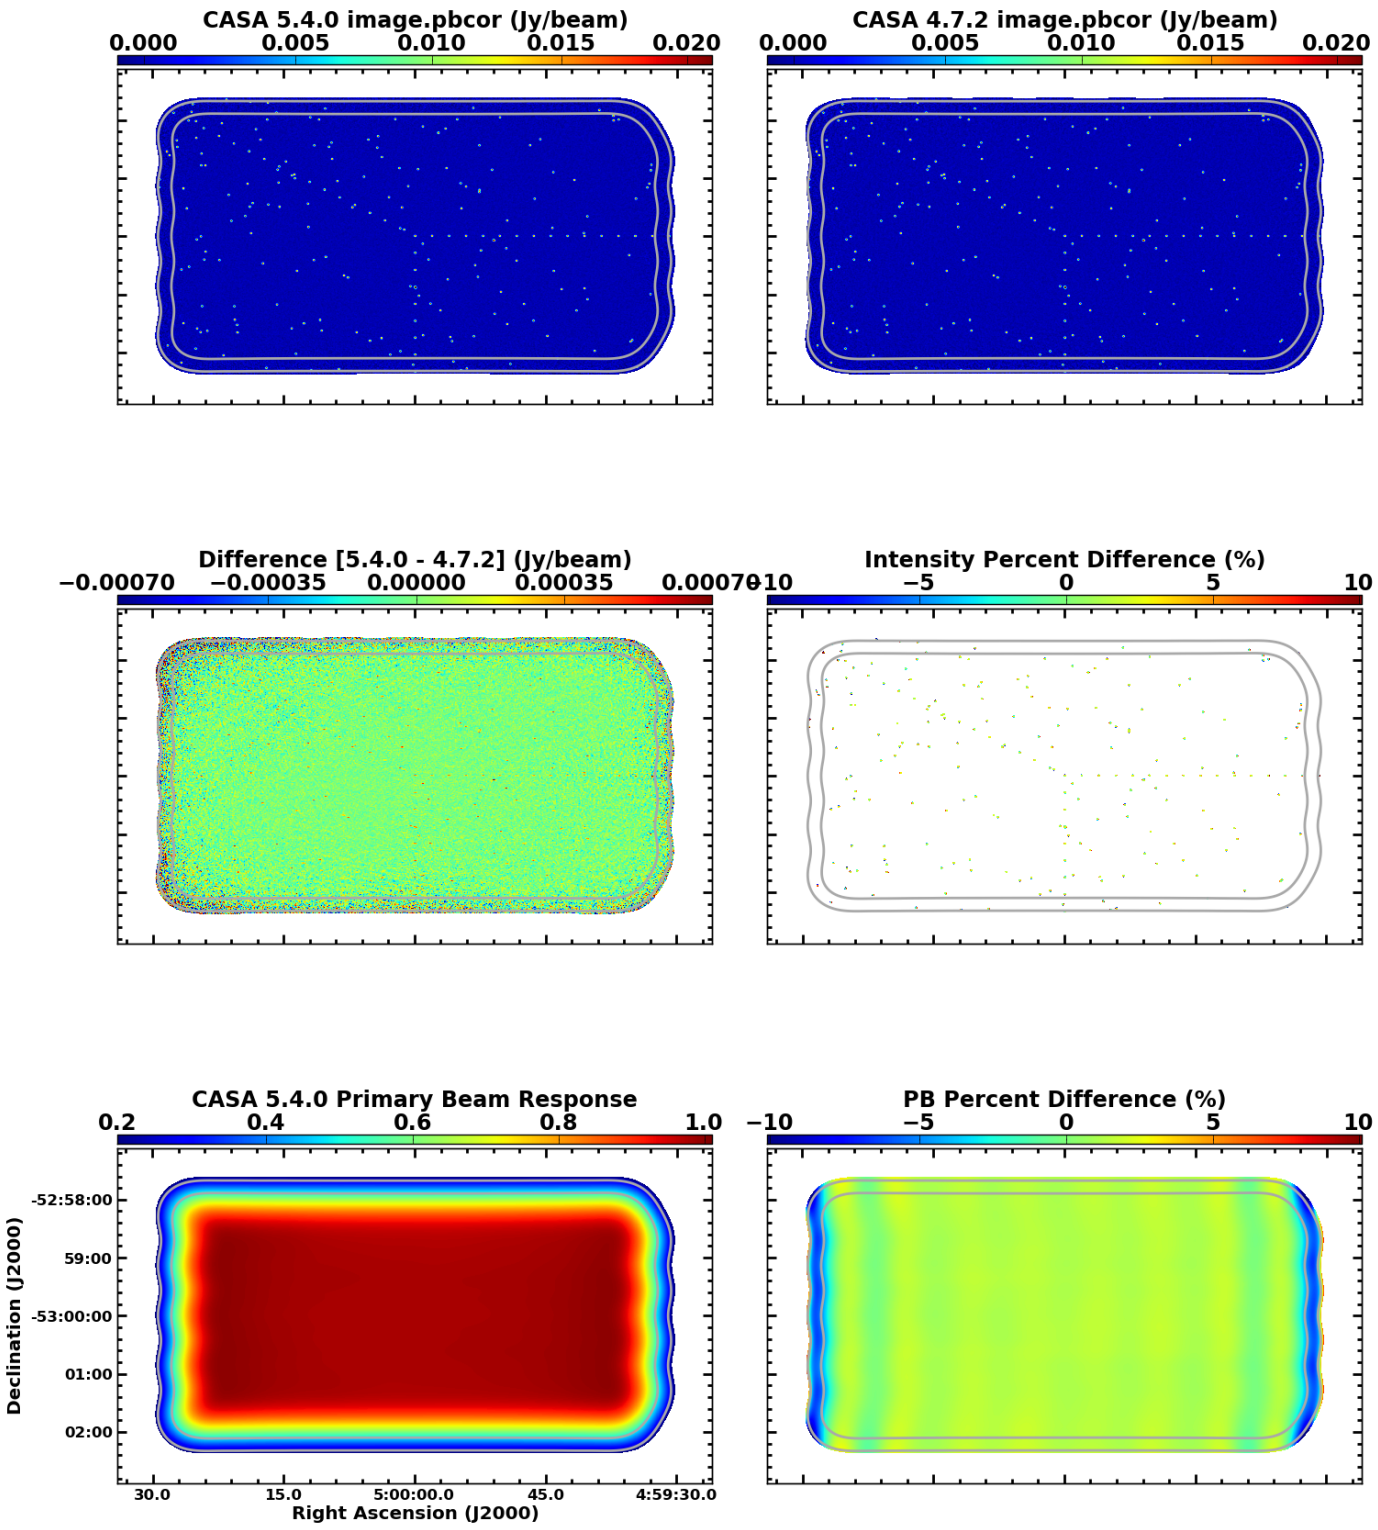

Figure 21: **Top row:** `tclean` images of a 12m array simulation with 16×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 4.7.2. Statistics of these images are given in row 12 of Tables 2 and 3.

B Additional examples of mosaic issue 2

B.1 7m 5.1.1 simulations

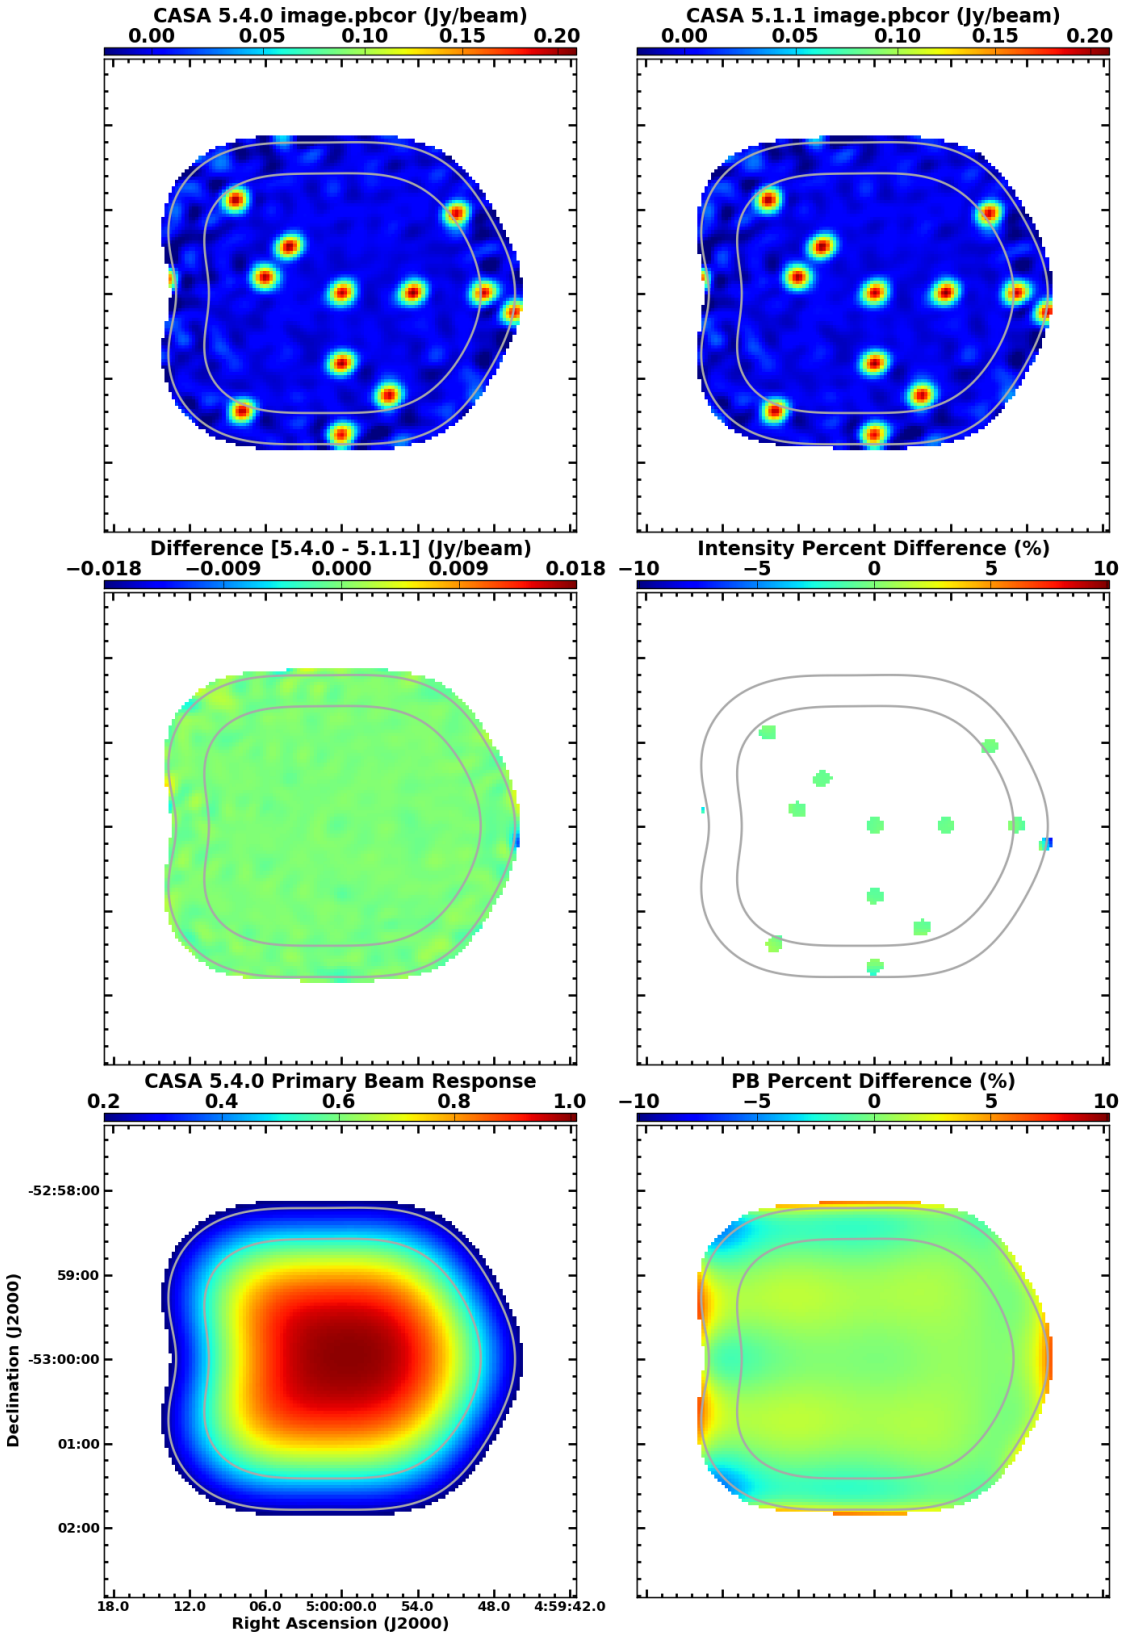

Figure 22: **Top row:** `tclean` images of a 7m array simulation with 3×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 1 of Tables 2 and 3.

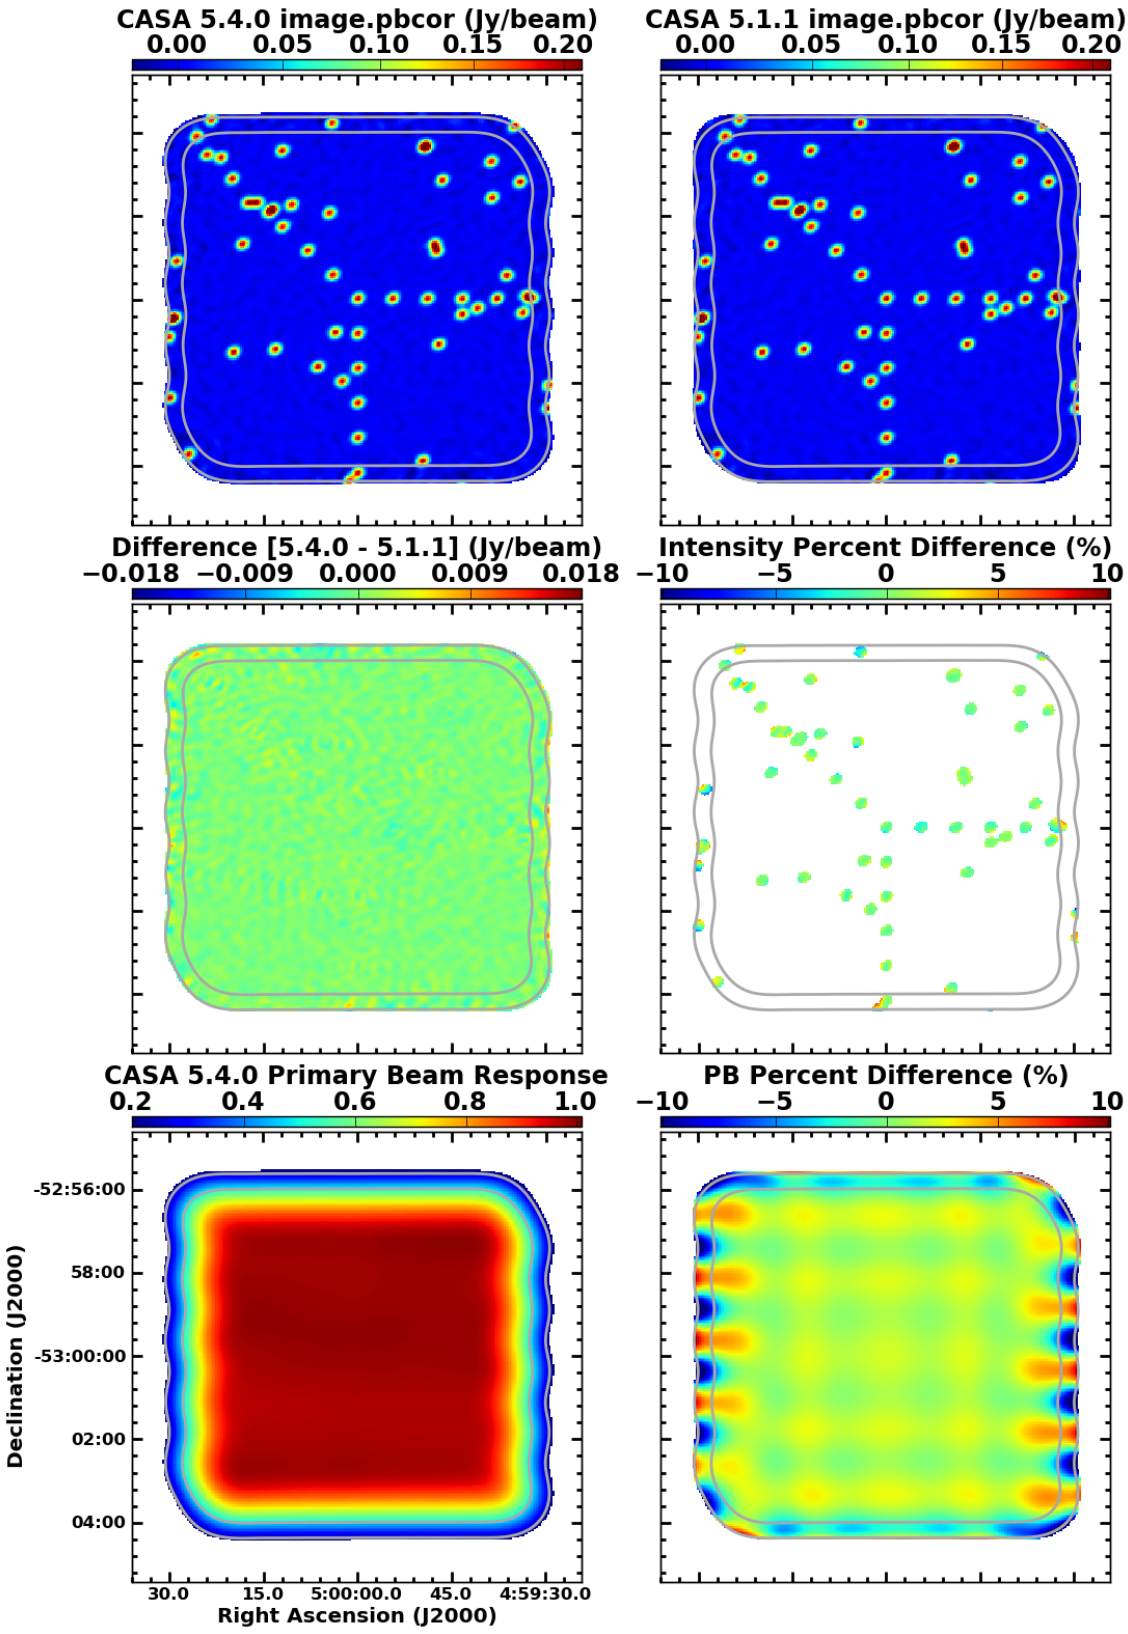

Figure 23: **Top row:** `tclean` images of a 7m array simulation with 9×10 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 3 of Tables 2 and 3.

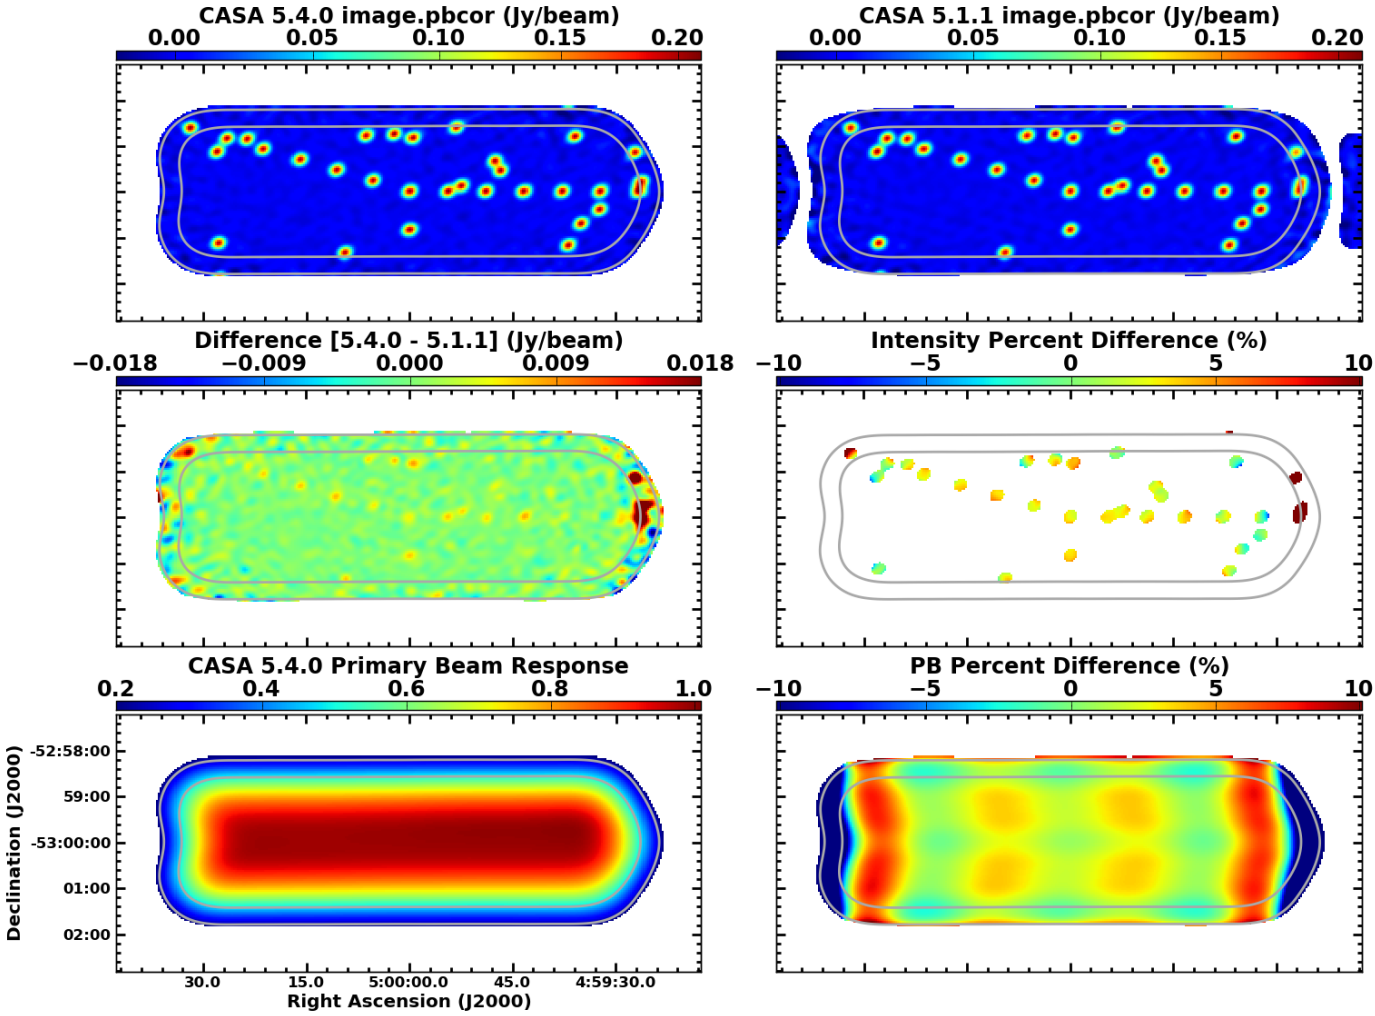

Figure 24: **Top row:** `tclean` images of a 7m array simulation with 11×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 4 of Tables 2 and 3.

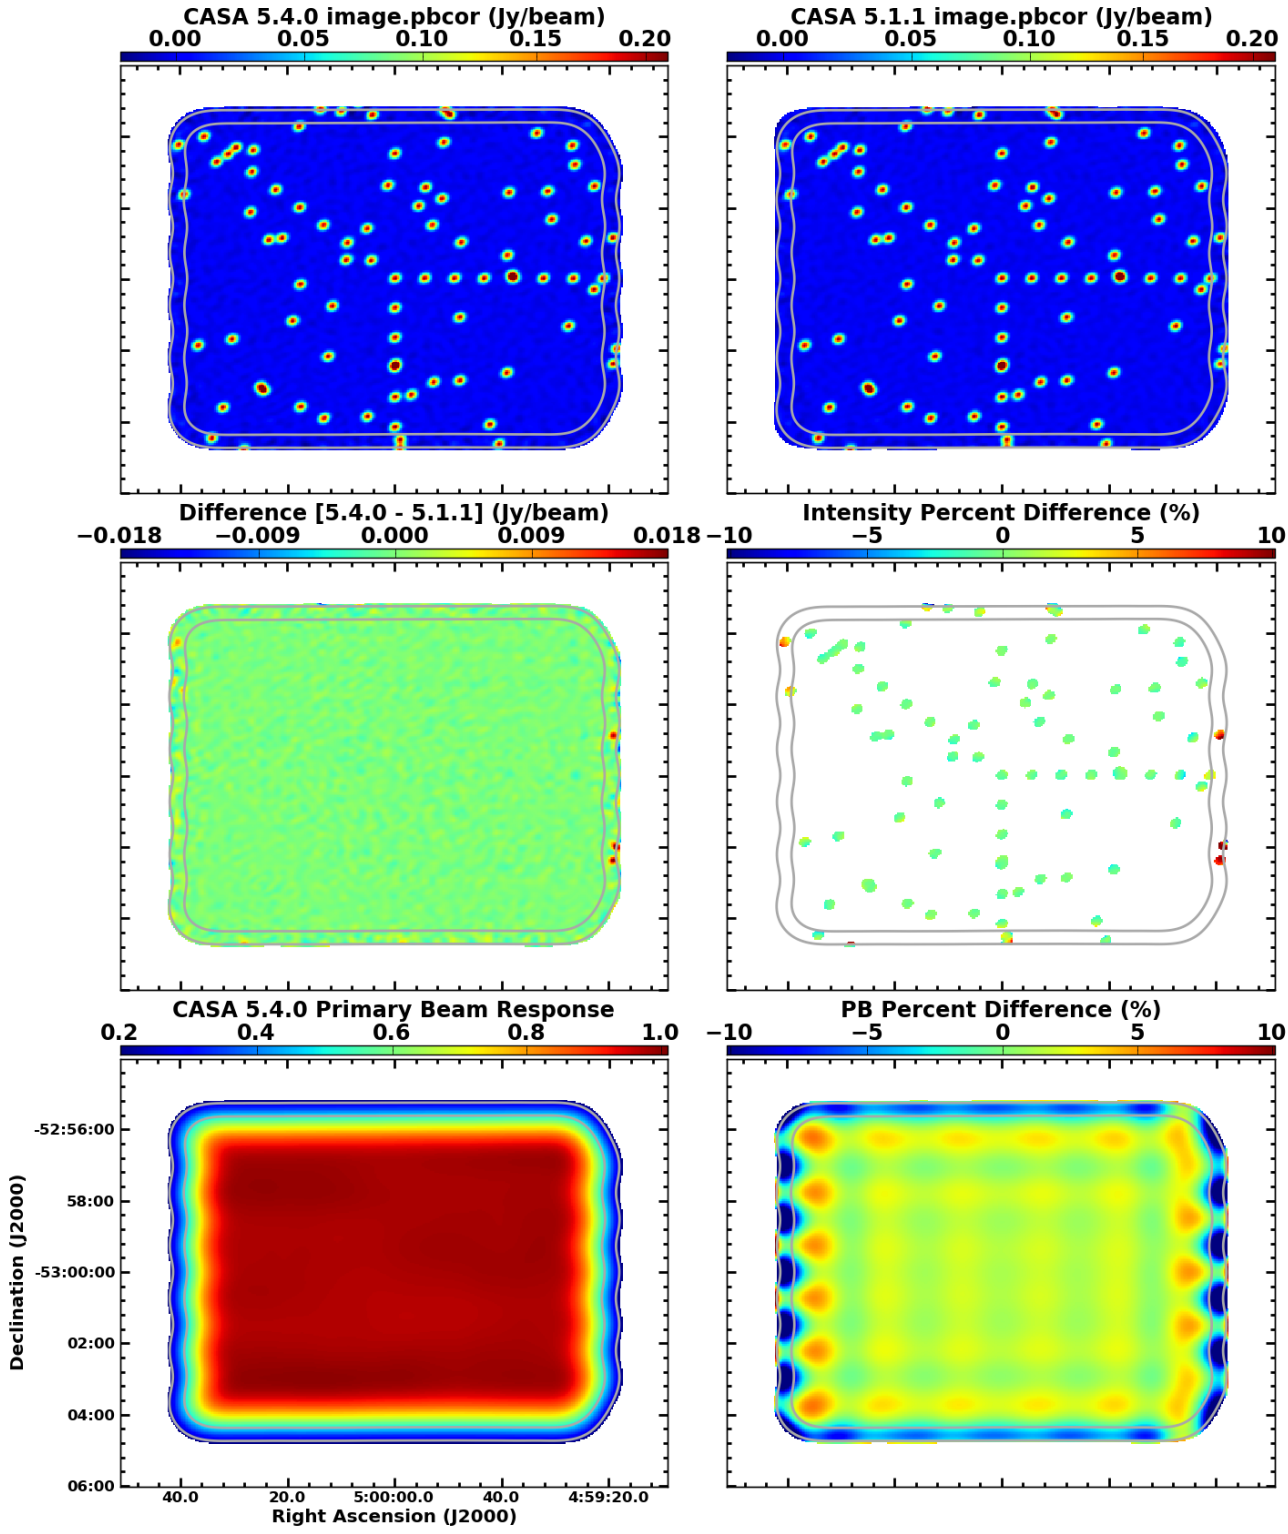

Figure 25: **Top row:** `tclean` images of a 7m array simulation with 13×11 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 5 of Tables 2 and 3.

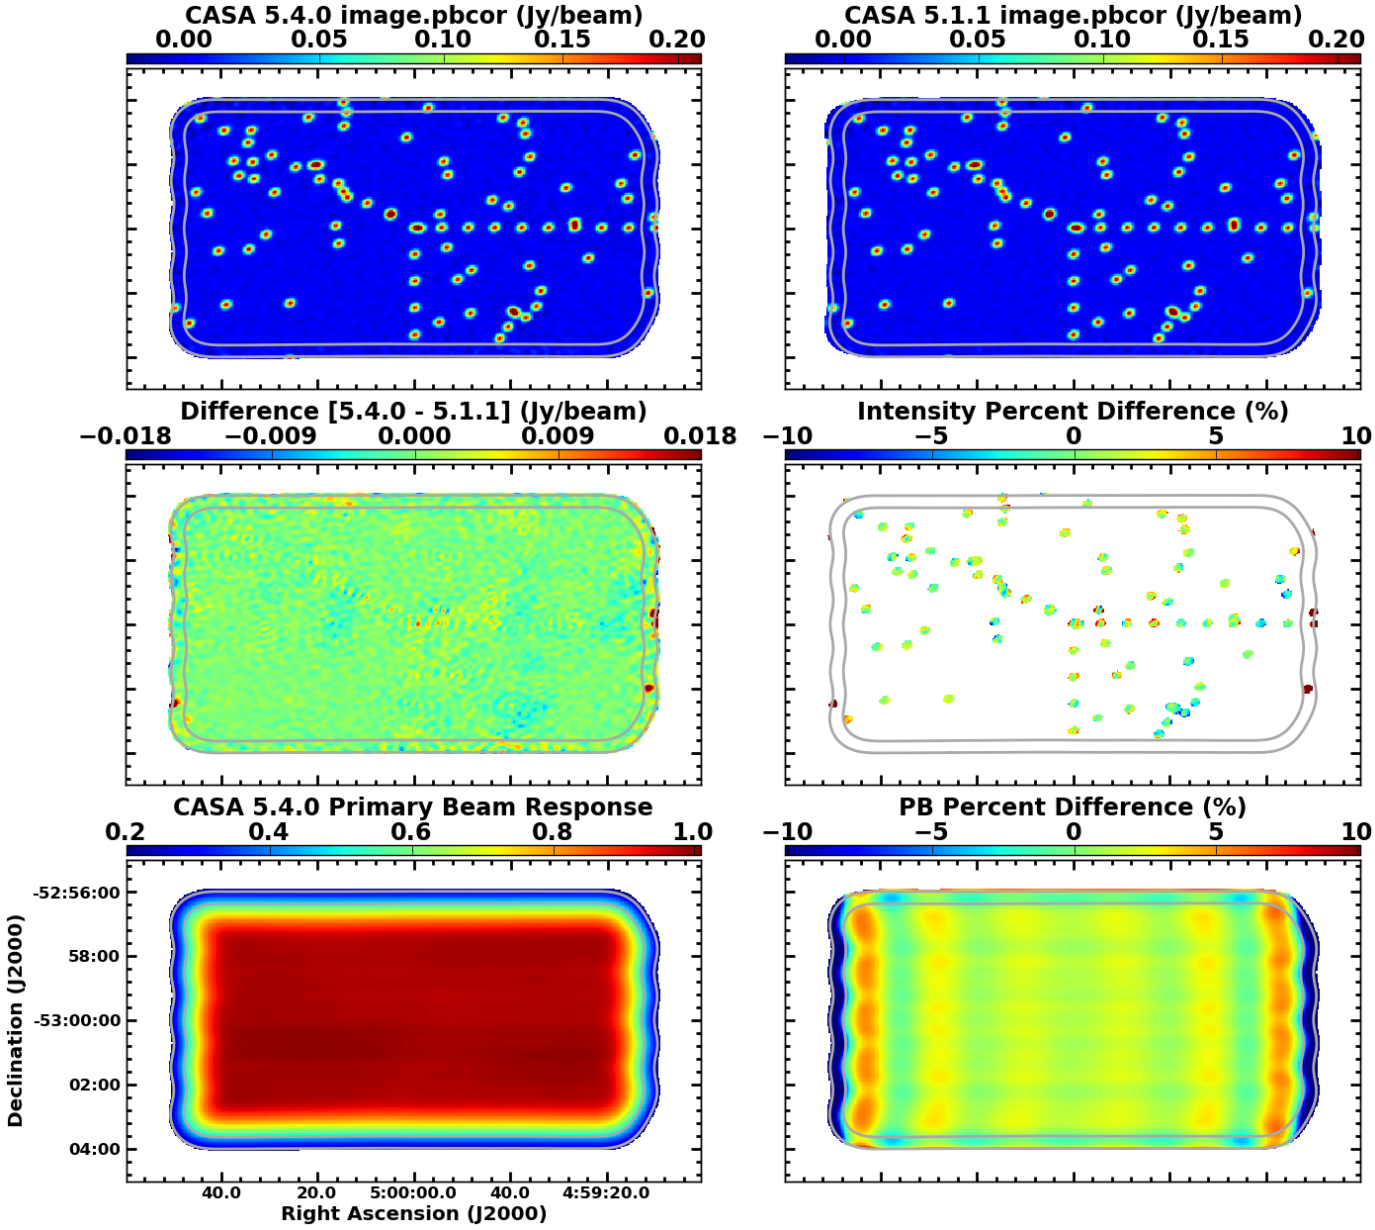

Figure 26: **Top row:** `tclean` images of a 7m array simulation with 16×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 6 of Tables 2 and 3.

B.2 12m 5.1.1 simulations

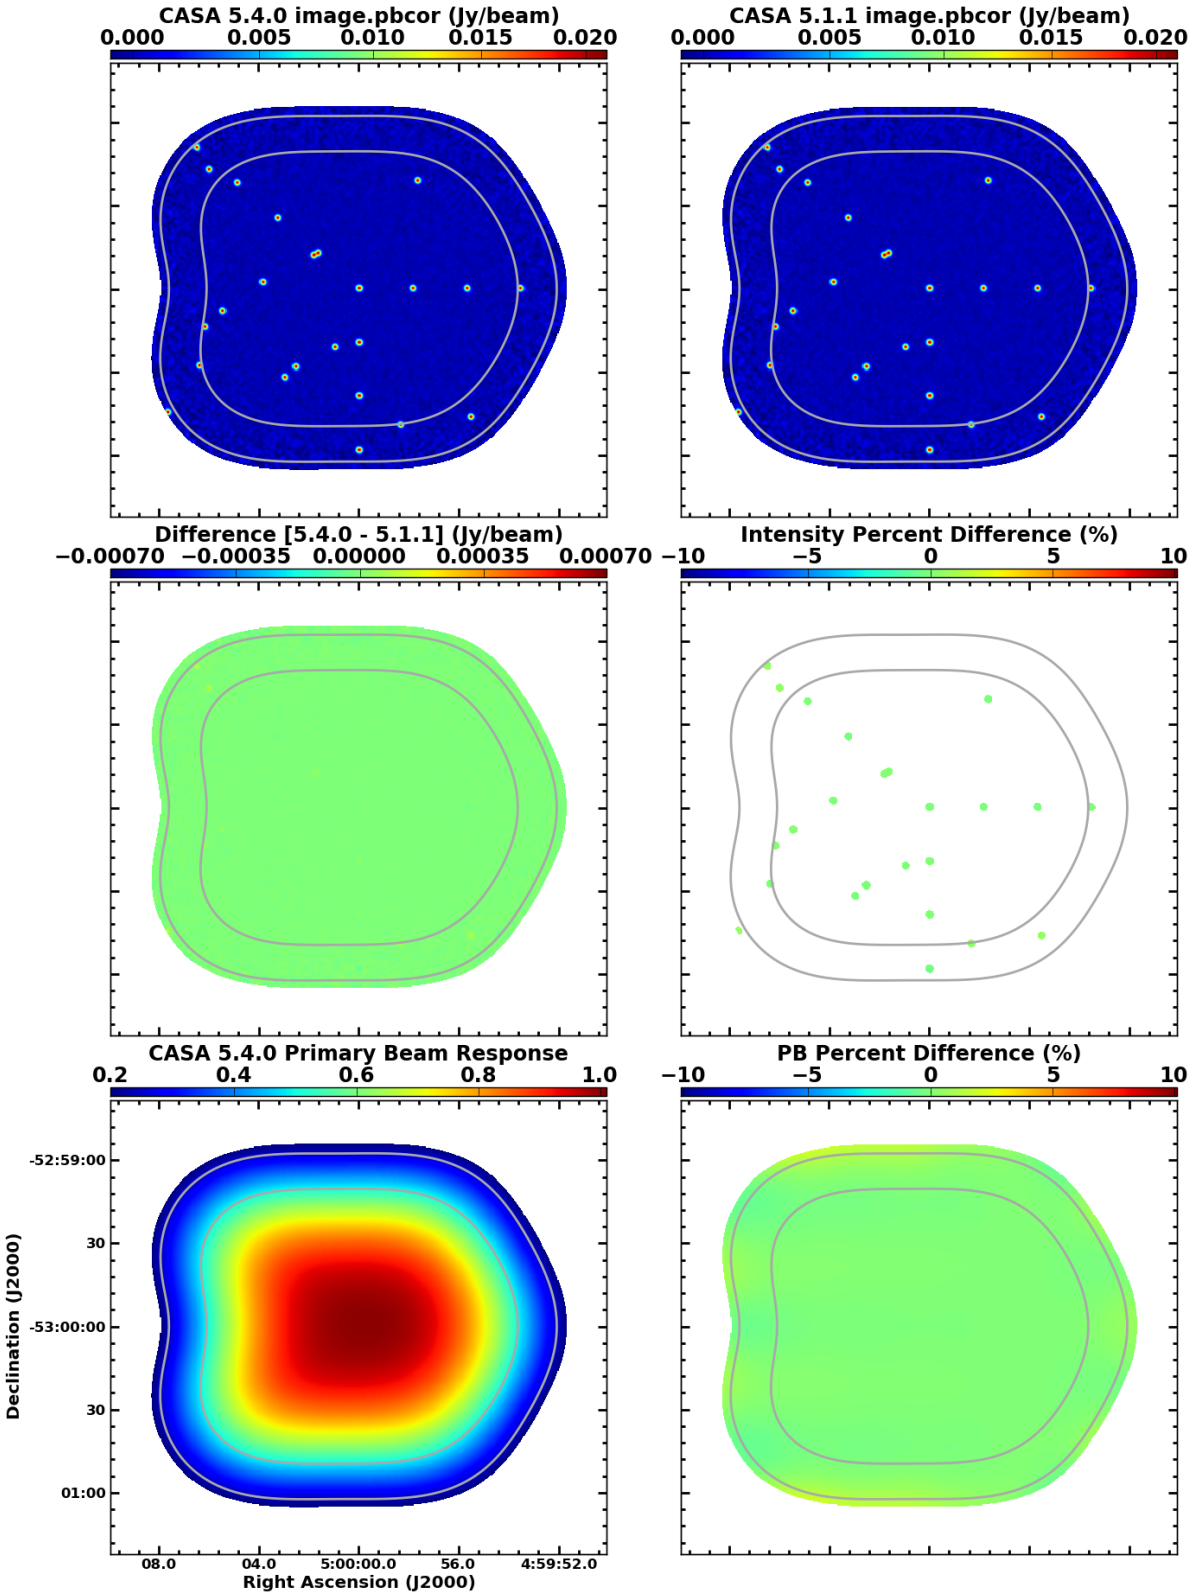

Figure 27: **Top row:** `tclean` images of a 12m array simulation with 3×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 7 of Tables 2 and 3.

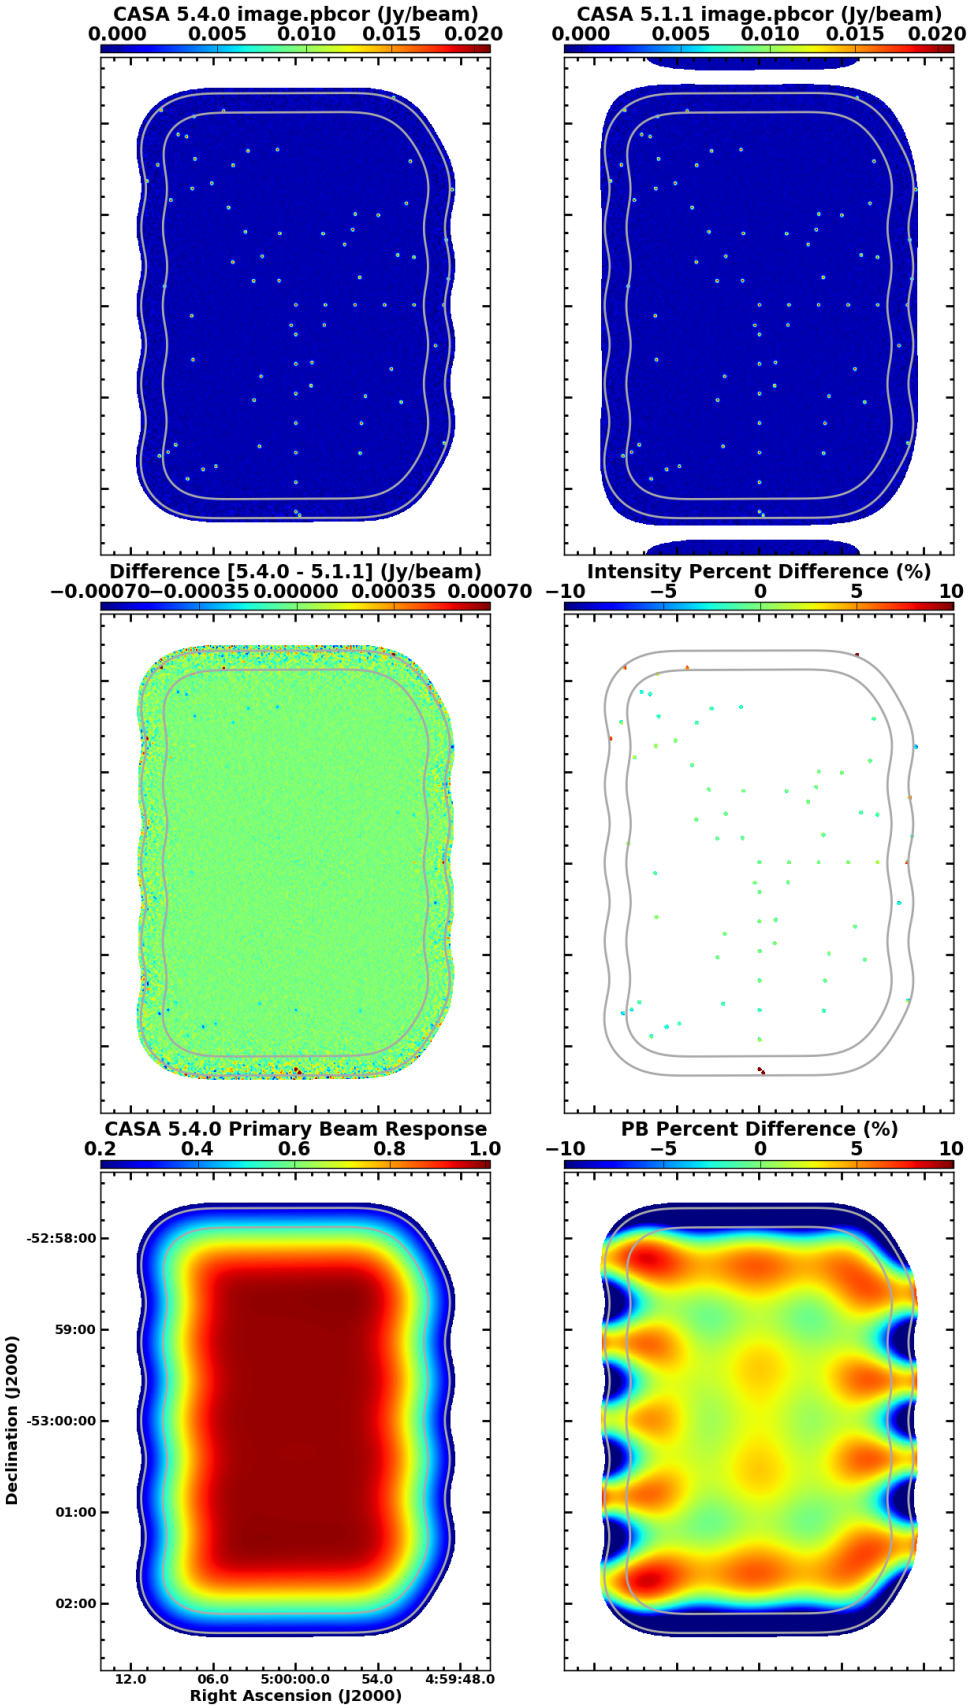

Figure 28: **Top row:** `tclean` images of a 12m array simulation with 5×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 8 of Tables 2 and 3.

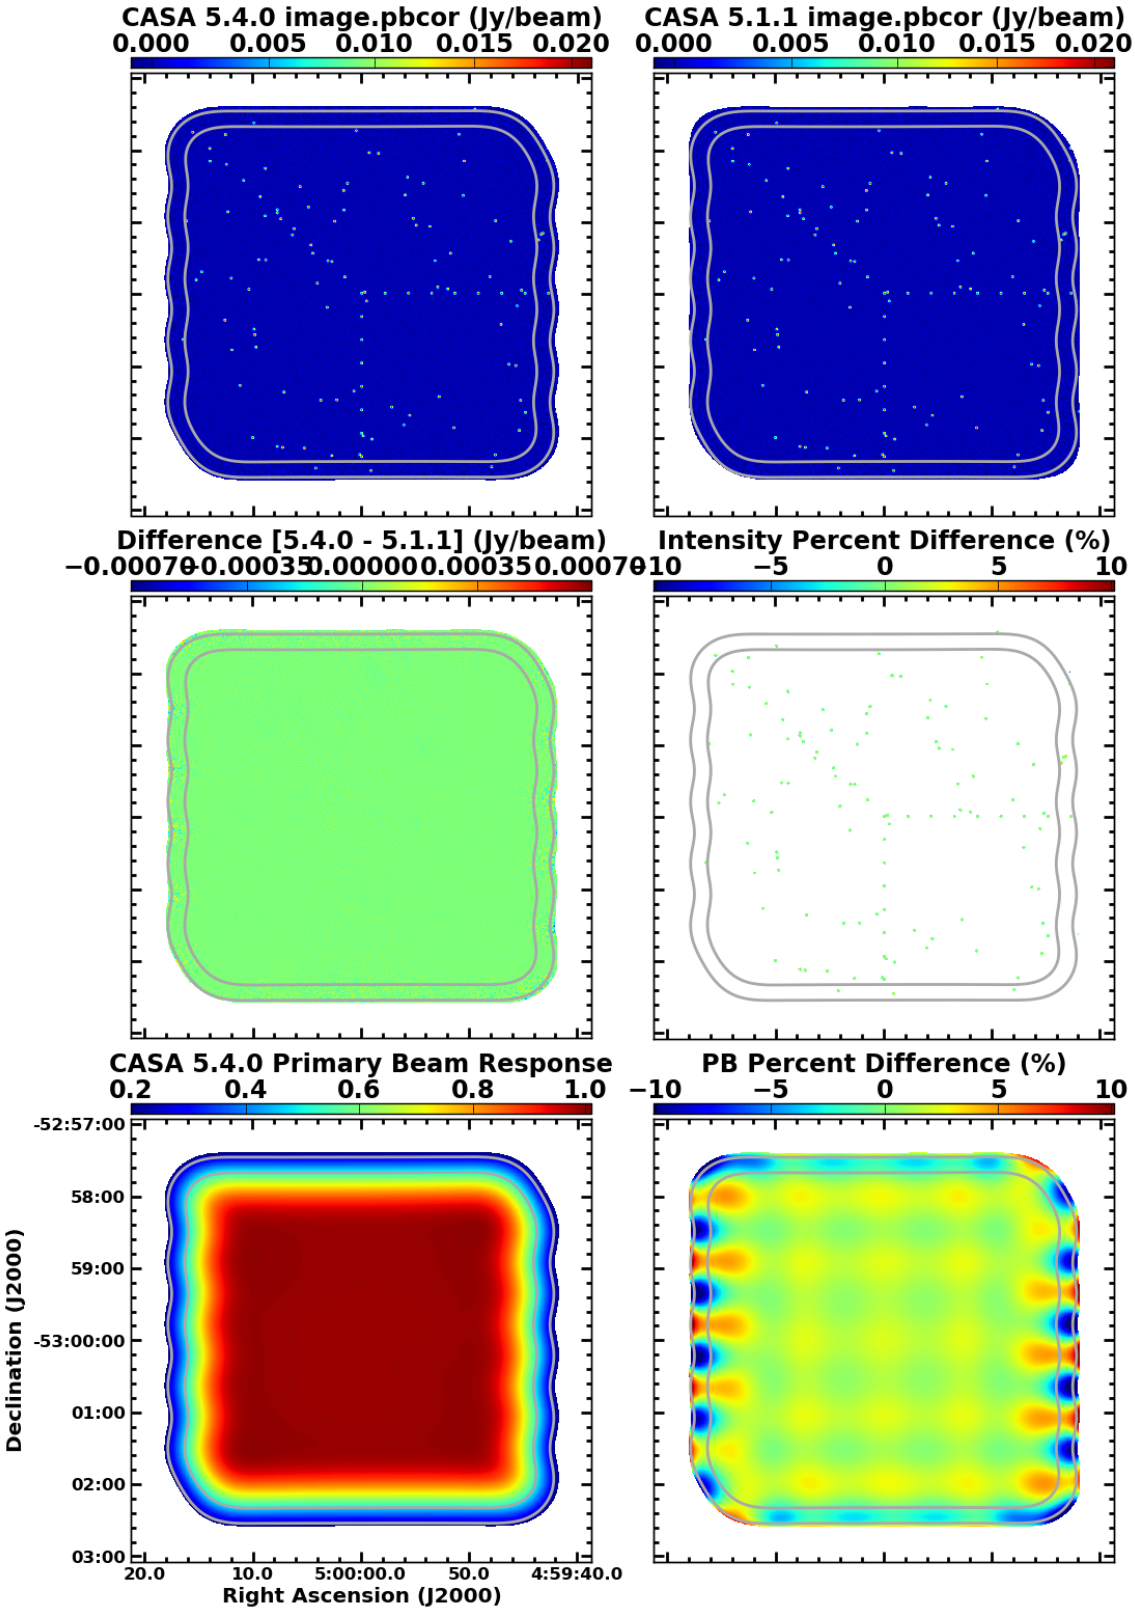

Figure 29: **Top row:** `tclean` images of a 12m array simulation with 9×10 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 9 of Tables 2 and 3.

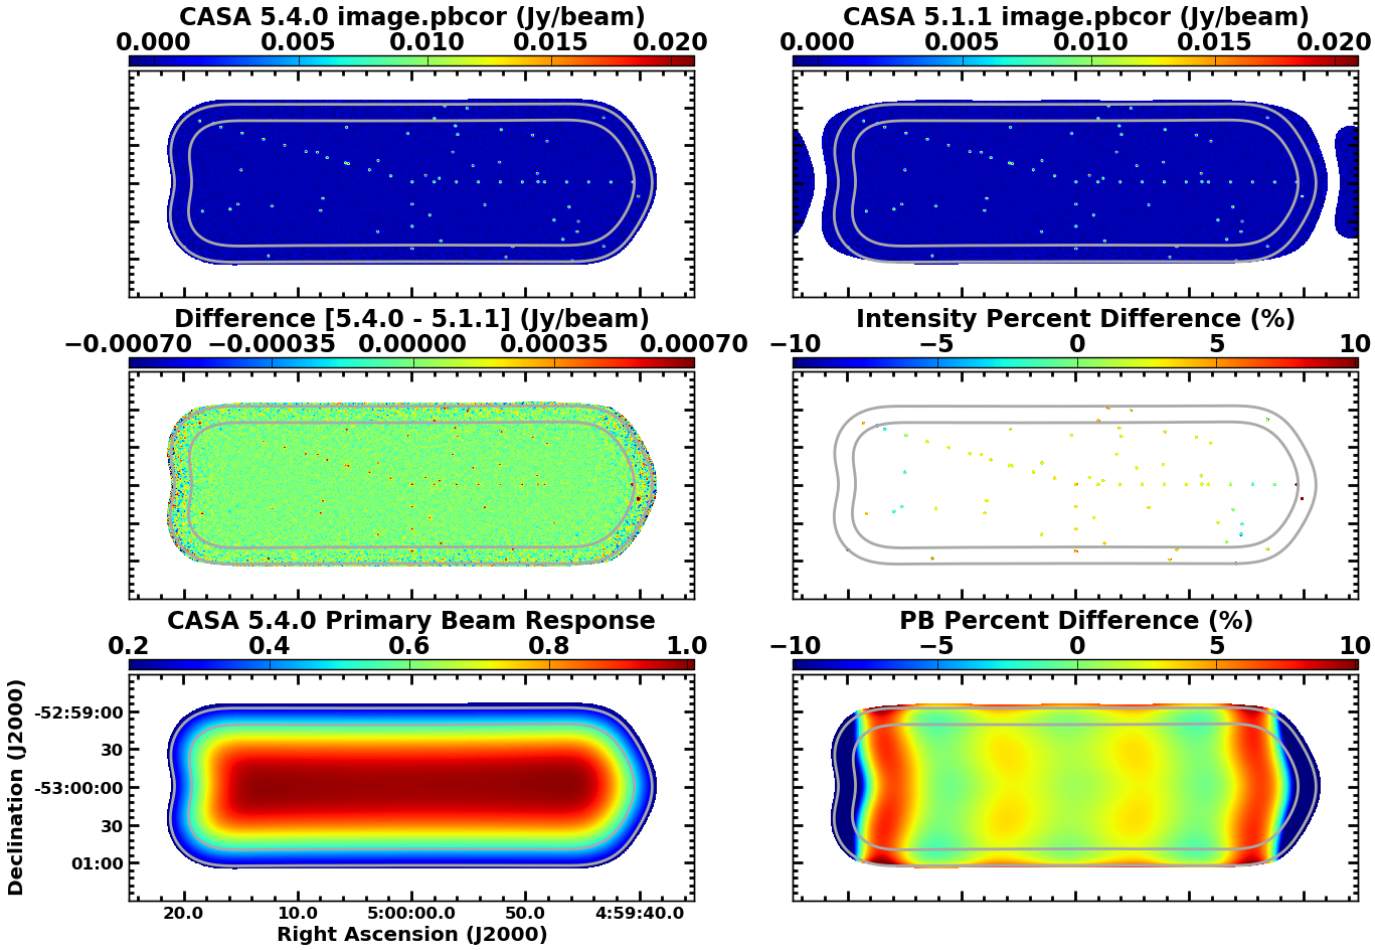

Figure 30: **Top row:** `tclean` images of a 12m array simulation with 11×3 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 10 of Tables 2 and 3.

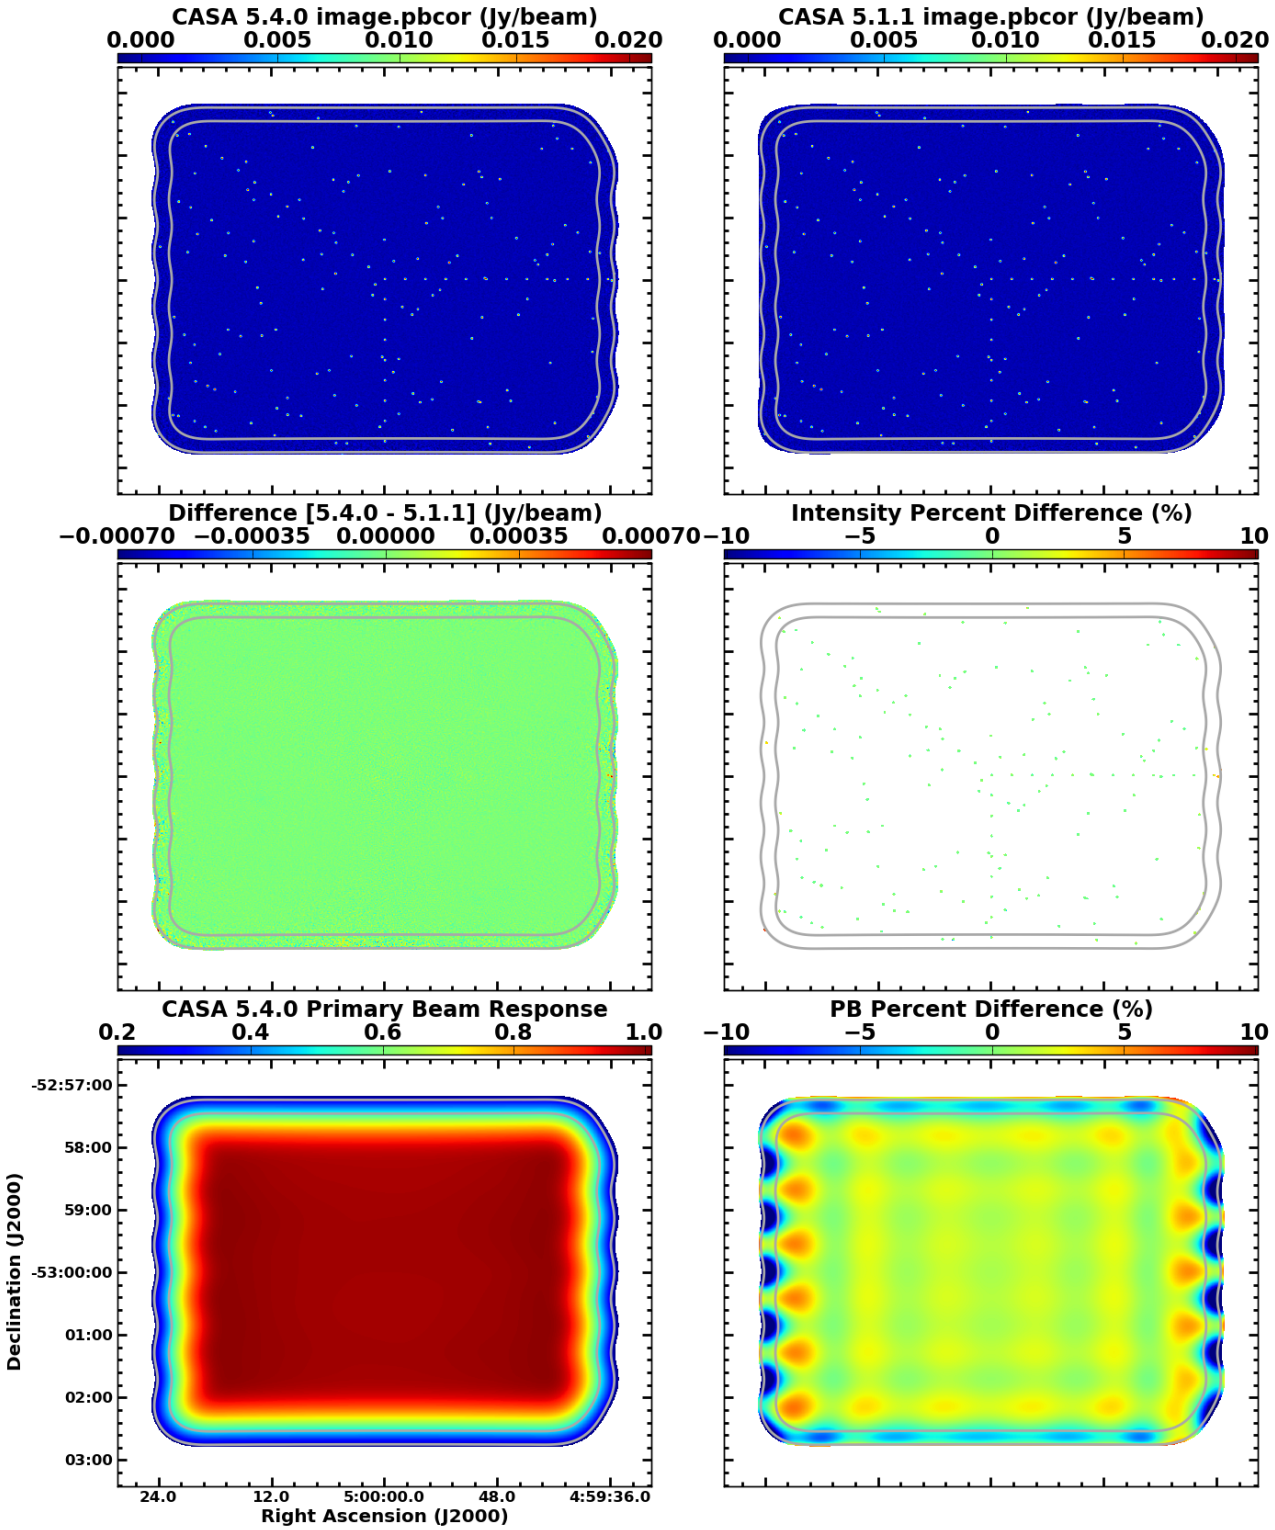

Figure 31: **Top row:** `tclean` images of a 12m array simulation with 13×11 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 11 of Tables 2 and 3.

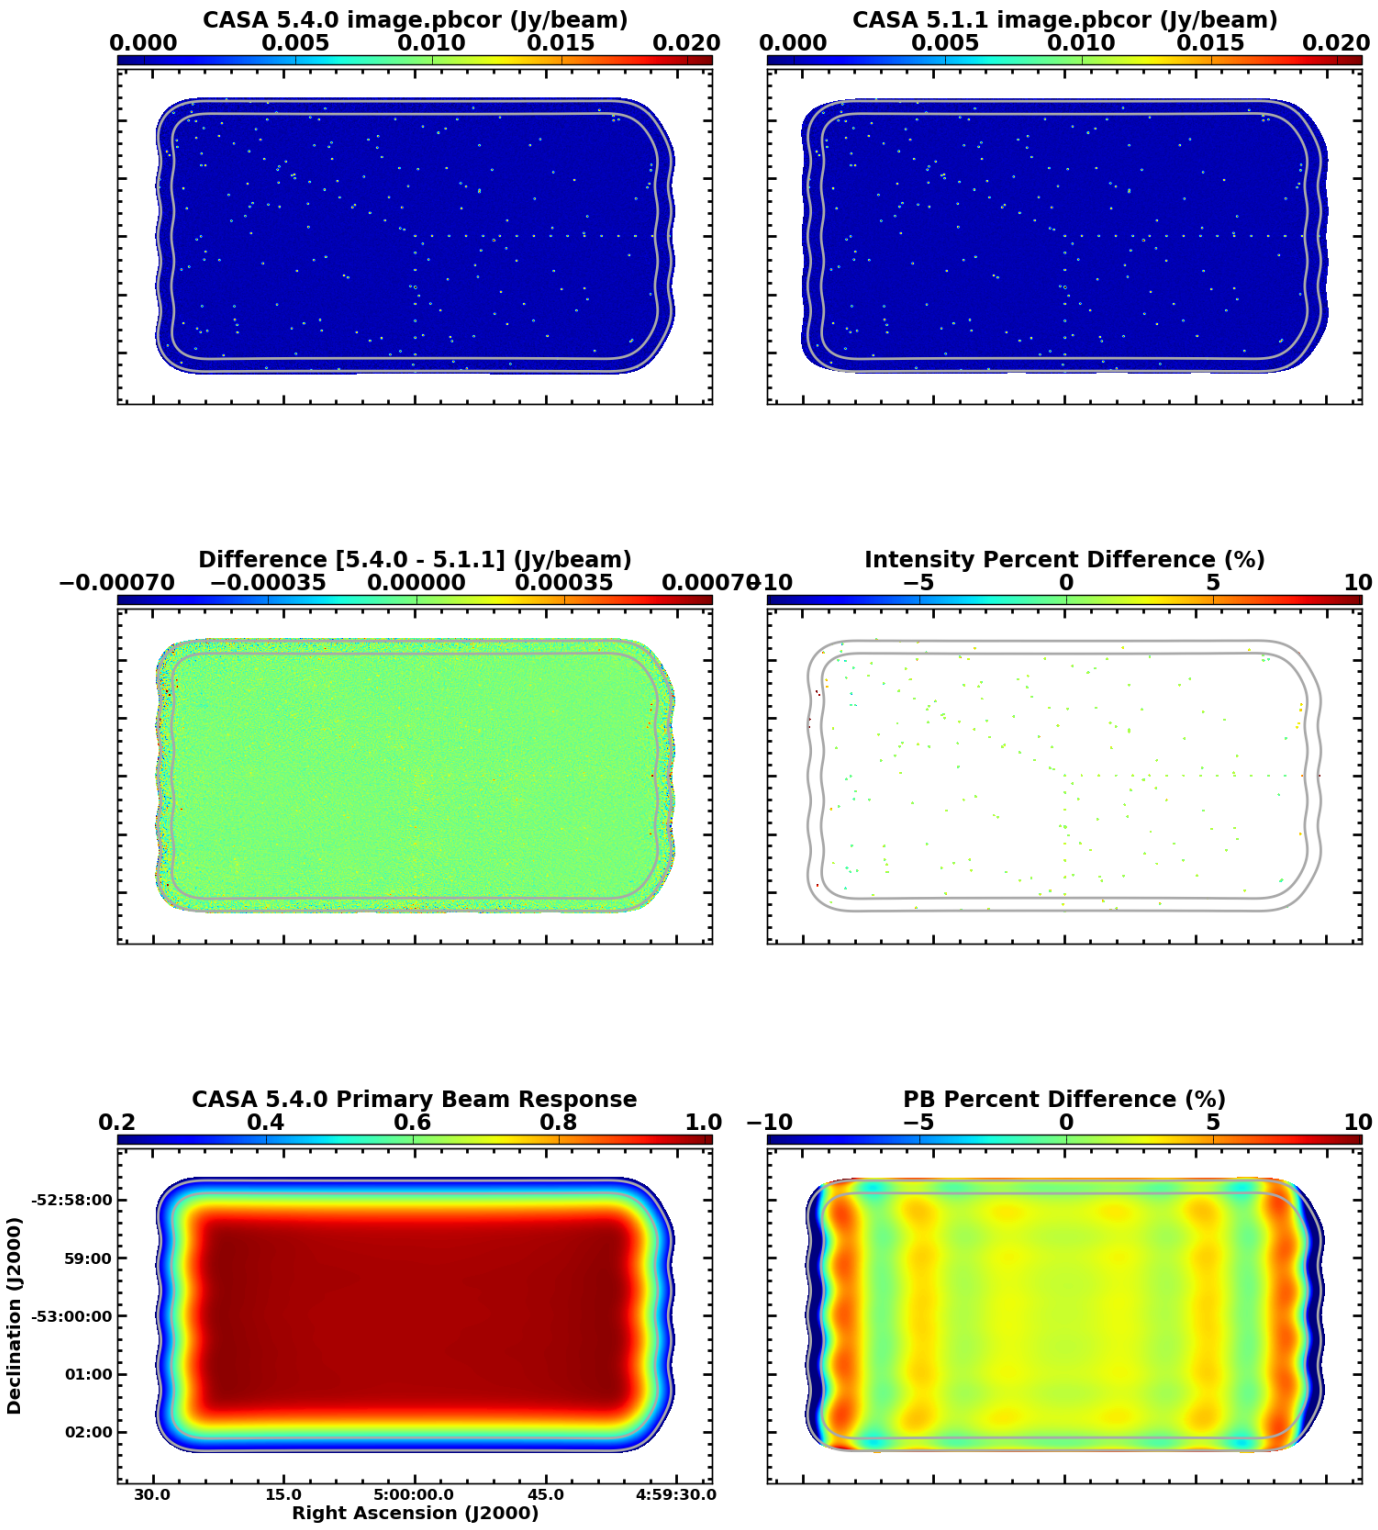

Figure 32: **Top row:** `tclean` images of a 12m array simulation with 16×9 pointings shown on the same intensity scale. **Middle row:** Difference of the images in the top row. The percentage image has been masked at 5σ. **Bottom row:** The primary beam (`.pb`) image from CASA 5.4.0 and the percentage difference from the corresponding `.pb` image from CASA 5.1.1. Statistics of these images are given in row 12 of Tables 2 and 3.