PHANGS–ALMA Data Processing and Pipeline

Adam K. Leroy^{1,40}, Annie Hughes^{2,3}, Daizhong Liu⁴, Jérôme Pety^{5,6}, Erik Rosolowsky⁷, Toshiki Saito⁴,
 Eva Schinnerer⁴, Andreas Schrubba⁸, Antonio Usero⁹, Christopher M. Faesi¹⁰, Cinthya N. Herrera⁵,
 Mélanie Chevance¹¹, Alexander P. S. Hygate¹², Amanda A. Kepley¹³, Eric W. Koch⁷, Miguel Querejeta⁹,
 Kazimierz Sliwa⁴, David Will¹, Christine D. Wilson¹⁴, Gagandeep S. Anand¹⁵, Ashley Barnes¹⁶, Francesco Belfiore¹⁷,
 Ivana Bešlić¹⁶, Frank Bigiel¹⁶, Guillermo A. Blanc^{18,19}, Alberto D. Bolatto²⁰, Mèdèric Boquien²¹, Yixian Cao²²,
 Rupali Chandar²³, Jérémy Chastenot²⁴, I-Da Chiang²⁴, Enrico Congiu^{18,19}, Daniel A. Dale²⁵, Sinan Deger²⁶,
 Jakob S. den Brok¹⁶, Cosima Eibensteiner¹⁶, Eric Emsellem^{27,28}, Axel García-Rodríguez⁹, Simon C. O. Glover²⁹,
 Kathryn Grasha³⁰, Brent Groves^{30,31}, Jonathan D. Henshaw⁴, María J. Jiménez Donaire⁹, Jaeyeon Kim¹¹,
 Ralf S. Klessen^{29,32}, Kathryn Kreckel¹¹, J. M. Diederik Kruijssen¹¹, Kirsten L. Larson²⁶, Janice C. Lee²⁶,
 Ness Mayker^{1,33}, Rebecca McElroy³⁴, Sharon E. Meidt³⁵, Angus Mok²³, Hsi-An Pan⁴, Johannes Puschner¹⁶,
 Alessandro Razza¹⁹, Patricia Sánchez-Bí'azquez^{36,37}, Karin M. Sandstrom²⁴, Francesco Santoro⁴, Amy Sardone^{1,33},
 Fabian Scheuermann¹¹, Jiayi Sun¹, David A. Thilker³⁸, Jordan A. Turner²⁵, Leonardo Ubeda³⁹, Dyas Utomo^{1,13,33},
 Elizabeth J. Watkins¹¹, and Thomas G. Williams⁴

¹ Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA; leroy.42@osu.edu

² CNRS, IRAP, 9 Av. du Colonel Roche, BP 44346, F-31028 Toulouse cedex 4, France

³ Université de Toulouse, UPS-OMP, IRAP, F-31028 Toulouse cedex 4, France

⁴ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117, Heidelberg, Germany

⁵ Institut de Radioastronomie Millimétrique (IRAM), 300 Rue de la Piscine, F-38406 Saint Martin d'Hères, France

⁶ Sorbonne Université, Observatoire de Paris, Université PSL, CNRS, LERMA, F-75014, Paris, France

⁷ Department of Physics, University of Alberta, Edmonton, AB T6G 2E1, Canada

⁸ Max-Planck-Institut für extraterrestrische Physik, Giessenbachstraße 1, D-85748 Garching, Germany

⁹ Observatorio Astronómico Nacional (IGN), C/Alfonso XII, 3, E-28014 Madrid, Spain

¹⁰ University of Massachusetts—Amherst, 710 N. Pleasant Street, Amherst, MA 01003, USA

¹¹ Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg, Mönchhofstraße 12-14, D-69120 Heidelberg, Germany

¹² Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

¹³ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903-2475, USA

¹⁴ Department of Physics and Astronomy, McMaster University, Hamilton, ON L8S 4M1, Canada

¹⁵ Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA

¹⁶ Argelander-Institut für Astronomie, Universität Bonn, Auf dem Hügel 71, D-53121 Bonn, Germany

¹⁷ INAF—Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50157, Firenze, Italy

¹⁸ Observatories of the Carnegie Institution for Science, 813 Santa Barbara Street, Pasadena, CA 91101, USA

¹⁹ Departamento de Astronomía, Universidad de Chile, Camino del Observatorio 1515, Las Condes, Santiago, Chile

²⁰ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

²¹ Centro de Astronomía (CITEVA), Universidad de Antofagasta, Avenida Angamos 601, Antofagasta, Chile

²² Aix Marseille Univ, CNRS, CNES, LAM (Laboratoire d'Astrophysique de Marseille), Marseille, France

²³ University of Toledo, 2801 W. Bancroft St., Mail Stop 111, Toledo, OH 43606, USA

²⁴ Center for Astrophysics and Space Sciences, Department of Physics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA

²⁵ Department of Physics and Astronomy, University of Wyoming, Laramie, WY 82071, USA

²⁶ Caltech-IPAC, 1200 E. California Blvd. Pasadena, CA 91125, USA

²⁷ European Southern Observatory, Karl-Schwarzschild Straße 2, D-85748 Garching bei München, Germany

²⁸ Univ Lyon, Univ Lyon 1, ENS de Lyon, CNRS, Centre de Recherche Astrophysique de Lyon UMR5574, F-69230 Saint-Genis-Laval, France

²⁹ Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Str 2, D-69120 Heidelberg, Germany

³⁰ Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia

³¹ International Centre for Radio Astronomy Research, University of Western Australia, 35 Stirling Highway, Crawley, WA 6009, Australia

³² Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Im Neuenheimer Feld 205, D-69120 Heidelberg, Germany

³³ Center for Cosmology and Astroparticle Physics, 191 West Woodruff Avenue, Columbus, OH 43210, USA

³⁴ Sydney Institute for Astronomy, School of Physics A28, The University of Sydney, NSW 2006, Australia

³⁵ Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281 S9, B-9000 Gent, Belgium

³⁶ Departamento de Física de la Tierra y Astrofísica, Universidad Complutense de Madrid, E-28040, Spain

³⁷ Instituto de Física de Partículas y del Cosmos IPARCOS, Facultad de Ciencias Físicas, Universidad Complutense de Madrid, E-28040, Spain

³⁸ Department of Physics and Astronomy, The Johns Hopkins University, Baltimore, MD 21218, USA

³⁹ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

Received 2020 September 26; revised 2021 February 2; accepted 2021 February 14; published 2021 July 19

Abstract

We describe the processing of the PHANGS–ALMA survey and present the PHANGS–ALMA pipeline, a public software package that processes calibrated interferometric and total power data into science-ready data products. PHANGS–ALMA is a large, high-resolution survey of CO(2–1) emission from nearby galaxies. The observations combine ALMA's main 12 m array, the 7 m array, and total power observations, and use mosaics of dozens to hundreds of individual pointings. We describe the processing of the $u-v$ data, imaging and deconvolution, linear

⁴⁰ This paper represents a collective effort by the PHANGS–ALMA data reduction team. Please see a description of contributions of individual team members in Appendix A.

mosaicking, combining interferometer and total power data, noise estimation, masking, data product creation, and quality assurance. Our pipeline has a general design and can also be applied to Very Large Array and ALMA observations of other spectral lines and continuum emission. We highlight our recipe for deconvolution of complex spectral line observations, which combines multiscale clean, single-scale clean, and automatic mask generation in a way that appears robust and effective. We also emphasize our two-track approach to masking and data product creation. We construct one set of “broadly masked” data products, which have high completeness but significant contamination by noise, and another set of “strictly masked” data products, which have high confidence but exclude faint, low signal-to-noise emission. Our quality assurance tests, supported by simulations, demonstrate that 12 m+7 m deconvolved data recover a total flux that is significantly closer to the total power flux than the 7 m deconvolved data alone. In the appendices, we measure the stability of the ALMA total power calibration in PHANGS–ALMA and test the performance of popular short-spacing correction algorithms.

Unified Astronomy Thesaurus concepts: [Radio interferometry \(1346\)](#); [Interstellar medium \(847\)](#); [CO line emission \(262\)](#); [Millimeter astronomy \(1061\)](#); [Astronomy data reduction \(1861\)](#); [Galaxies \(573\)](#); [Star formation \(1569\)](#)

1. Introduction

Modern radio interferometric data sets often include hundreds or thousands of distinct observations. They combine data from different arrays, including both total power and interferometric measurements, and both the visibility and image data have large volumes. Calibrating, imaging, and deconvolving these data to produce correct images of the sky can be challenging. Even after these steps, further processing is required to translate these images (or data cubes) into data products ready for scientific analysis.

Current observatories, especially the Atacama Large Millimeter/submillimeter Array (ALMA), have made amazing strides toward automated, high-quality calibration of interferometric data. For ALMA, this stems from hard work by the observatory and the success of the ALMA interferometric and total power pipelines (L. Davis et al. 2021, in preparation), which in turn build on the CASA software project (Common Astronomy Software Applications; McMullin et al. 2007). Thanks to these efforts, ALMA delivers well-calibrated visibility ($u-v$) data to its users.

Turning these calibrated $u-v$ data into science-ready data products represents a complex task. In this paper, we focus on ALMA observations of CO line emission from nearby galaxies. This emission has complex spatial and velocity structure. It often spans across many individual telescope pointings, and requires both high angular resolution and short-spacing data to recover a full picture of the emission. Moreover, most scientific analysis does not make use of the full position–position–velocity data cube produced by imaging. Translating the visibility data into a science-ready form also involves producing a suite of higher-level data products with well-understood properties and uncertainties.

This paper describes the post-calibration processing pipeline constructed to carry out these steps for PHANGS–ALMA, a CO survey of nearby galaxies. As part of this project, we encountered all the issues mentioned: large data volume, the need to reconstruct complex emission from observations using multiple arrays and telescopes, and the need to create high-level data products for use in scientific analysis. We address them by adopting or developing appropriate algorithms and implementing them in a modular python and CASA pipeline. The result, described in this paper, is a suite of reproducible, automated methods for processing calibrated $u-v$ observations of galaxies into science-ready data products.

1.1. PHANGS–ALMA

PHANGS–ALMA is an ALMA survey of CO(2–1) emission from 90 nearby galaxies. The sample selection, observations,

and scientific motivation are described by Leroy et al. (2021). In brief, this is a large, multicycle program focused on mapping CO(2–1) emission at ~ 100 pc resolution across the areas of active star formation in a large, cleanly selected sample of nearby galaxies. The core of the survey is an ALMA Cycle 5 Large Program (P.I. E. Schinnerer), which is supplemented by a series of smaller programs across five ALMA observing cycles.

Tables 1 and 2 summarize the properties of the PHANGS–ALMA data set. The survey combines observations with ALMA’s main 12 m array and both parts of the Morita Atacama Compact Array (ACA): the 7 m array and the total power antennas. The 12 m array observations used relatively compact configurations, corresponding to angular resolutions of $\sim 1''$ – $1''.5$ at the frequency of CO(2–1). ALMA’s main array and 7 m array observe independently, so that separate $u-v$ data exist for the main 12 m array and the array of smaller 7 m antennas. The 7 m array consists of fewer antennas, 12 in total. As a result, the total integration time needed to achieve suitable surface brightness sensitivity using the 7 m array is 3–7 times longer than that of the main array (see ALMA Technical Handbook).⁴¹

We covered each target using large, multifield mosaics with sizes that frequently approach the observatory-imposed maximum of 150 pointings. When one 150-field mosaic could not cover the galaxy, we observed multiple, adjacent mosaics to cover the galaxy. The correlator setup covered $^{12}\text{CO}(2-1)$ at high spectral resolution and one or more other lines at coarser spectral resolution. We devoted the remainder of the correlator resources to observe the continuum.

In more basic terms, the data for each PHANGS–ALMA target consist of single-dish spectroscopic mapping and interferometric visibilities, or “ $u-v$ data,” for dozens or hundreds of individual pointed fields. The 7 m and 12 m arrays map almost the same area on the sky, but do not share the same pointing centers. The total power data consist of individual spectra obtained using on-the-fly mapping techniques that cover the same spatial region mapped by the interferometer.

Based on the inspection described in Section 3, we verified that, as expected, ALMA delivers reliable, well-calibrated $u-v$ data. These data products reflect the excellent performance of the ALMA interferometric calibration pipeline, the stability of the instrument, and the still-minimal impact of radio frequency interference (RFI) on millimeter-wave observations.

⁴¹ For Cycle 5, when the large program was executed: [https://almascience.nrao.edu/documents-and-tools/\[0\]cycle5/\[0\]alma-technical-handbook](https://almascience.nrao.edu/documents-and-tools/[0]cycle5/[0]alma-technical-handbook).

Table 1
PHANGS–ALMA Data Summary

Description	Value
Galaxies	90
Targets (i.e., individual mosaics)	136
Galaxies observed using multiple mosaics	26
ACA 7 m array measurement sets	479
Typical 7 m $u-v$ range	8–43 m (6–33 k λ)
Typical 7 m beam	7".2 $\times$ 4".4
Main 12 m array measurement sets	184
Typical 12 m $u-v$ range	13–380 m (10–292 k λ)
Typical 12 m+7 m beam	1".26 $\times$ 1".04
Total power observations	744
Typical total power beam	28".4
CO(2–1) native channel	~ 0.32 km s $^{-1}$
CO(2–1) working channel ^a	~ 2.54 km s $^{-1}$
C ¹⁸ O(2–1) native channel	~ 2.7 km s $^{-1}$
C ¹⁸ O(2–1) working channel ^a	~ 6.0 km s $^{-1}$
Total bandwidth	6.7 GHz

Notes. These numbers refer to the data processed by our team just prior to the first public data release of PHANGS–ALMA. For details of the public data release and a summary of the survey and observations see Leroy et al. (2021).

^a Target velocity resolution. The pipeline aims to get as close as possible to this number without going over.

1.2. From $u-v$ Data to Science-ready Data Products

While calibration is handled by the observatory, the observatory does not deliver images that combine multiple arrays, interferometric and total power data, or multiple mosaics. Nor does the observatory currently provide derived data products beyond data cubes and images. This leaves the user with the task of translating the visibility and total power data into science-ready data products.

This procedure begins with imaging and deconvolution. The $u-v$ data sample the Fourier transform of the sky emission at each frequency. They need to be gridded and Fourier transformed, or “imaged,” at each frequency to produce data cubes. Interferometers sample the $u-v$ plane incompletely. Producing accurate images of the sky requires reconstructing the true intensity distribution from these incomplete visibility data. This process is referred to as deconvolution—or often simply as “CLEANing,” in reference to the most commonly used algorithm (Högbom 1974). Modern methods include both versions of the classic CLEAN (Högbom 1974), which reconstructs the emission as a collection of point sources, and the more recent “multiscale CLEAN” (Cornwell 2008), which uses a combination of Gaussian components with a range of scales to reconstruct the image. In parallel, the total power data need to have any frequency-dependent baseline structure removed and the data then combined from a spatially sampled grid of individual spectra into data cubes (e.g., see Mangum et al. 2007).

After deconvolution, the interferometric data need to be combined with the single-dish data in order to correct for the interferometer’s lack of sensitivity to extended emission. Approaches to this step vary, and include joint imaging of the interferometric and total power data (e.g., Koda et al. 2019), image plane combination (Stanimirovic et al. 1999), or Fourier-based processing (“feathering”; Cotton 2017). For galaxies observed, imaged, and deconvolved in separate parts, the individual parts must also be stitched together after imaging.

We use linear mosaicking to combine individual parts and yield a complete image of each galaxy.

The steps described above yield science-ready data cubes. The subsequent analysis often relies on higher-level data products—for example, maps of integrated line intensity, mean velocity, or line width, as well as more complex quantities. The first step toward creating such high-level products is usually signal identification. For line emission from well-resolved galaxies, the fraction of a data cube filled by real emission is often small, reflecting the wide bandwidth of the instrument compared to typical line widths for the interstellar medium (ISM). Identifying the parts of the data cube likely to contain emission is critical to accurately measuring the moments of the emission distribution, particularly the higher moments like line width (“moment 2”).

The most common approach to signal identification is to “mask” the data cubes. In this procedure each voxel, i.e., each three-dimensional volumetric pixel, is labeled “True” or “False” according to whether it is likely to contain line emission (“True”) or only noise (“False”). Choices made during the masking process can prioritize either high completeness, meaning inclusion of all emission, or a low false-positive rate, meaning that “True” pixels are very likely to contain real emission.

After identifying the part of a cube likely to contain signal, the mask is applied to the line data cube. The voxels containing signal are then “collapsed” to form maps that describe the line emission in ways directly relevant to scientific analysis. The resulting maps are usually referred to as “moment” maps, though this term frequently includes more than just the intensity-weighted velocity moments of the data. Commonly computed quantities include line-integrated intensity, measurements of the line width and spectral profile shape, and measurements of the characteristic velocity.

1.3. The PHANGS–ALMA Pipeline

The ALMA imaging interferometric pipeline implements deconvolution and imaging of visibility data for individual arrays (e.g., Kepley et al. 2020), but does not yet image combined data from different arrays or combine total power and interferometric data. These steps are all necessary to produce science-ready data products to achieve our science goals. This paper describes the steps taken to postprocess the PHANGS–ALMA data and details the motivations for our choices. We also describe the PHANGS–ALMA postprocessing pipeline software, which combines CASA with Python extended by additional packages (Section 2.3).

Although the PHANGS pipeline was developed for the PHANGS–ALMA survey to produce CO(2–1), C¹⁸O(2–1), and continuum images, the software represents a general post-processing pipeline. We have used it to process CO(3–2), CO(4–3), ¹³CO(2–1), HCN(1–0), HCO⁺(1–0), CS(2–1), [C I](³P_{1–3}P₀), and dust continuum data from ALMA as well as HI 21 cm data from the Very Large Array (VLA). Altogether, we have processed on the order of 1000 interferometric observations using this software. The closely related total power calibration and imaging pipeline presented in Herrera et al. (2020) and summarized here has also processed on the order of 1000 total power observations.

Section 2 summarizes the workflow, notes the software used to implement the PHANGS pipeline, and defines key terms. Sections 3 and 4 describe the $u-v$ data processing, imaging,

Table 2
PHANGS–ALMA CO(2–1) Imaging

Description	ACA 7 m Value (minimum — 16th percentile — median — 84th percentile — maximum)	12 m+7 m Value (minimum — 16th percentile — median — 84th percentile — maximum)
Beam		
Major axis [″]	6.2 — 6.8 — 7.2 — 7.9 — 9.7	0.58 — 1.0 — 1.2 — 1.6 — 1.9
Position angle [°]	69 — 82 — 88 — 98 — 124	5 — 59 — 95 — 116 — 179
Elongation [major/minor axis]	1.1 — 1.4 — 1.7 — 2.0 — 2.3	1.0 — 1.1 — 1.2 — 1.4 — 1.9
Pixels across beam minor axis	3.5 — 3.8 — 4.4 — 4.9 — 5.9	4.1 — 4.9 — 5.9 — 7.0 — 8.7
Area imaged^a		
Pixels across cube major axis	120 — 240 — 288 — 384 — 512	720 — 1152 — 1536 — 2304 — 4608
Area mapped [arcmin ²]	1.2 — 4.0 — 8.2 — 22.2 — 22.8	0.7 — 2.8 — 6.2 — 7.8 — 15.2
Spatial dynamic range [$\sqrt{\text{area}/\text{beam}}$]	11 — 22 — 28 — 42 — 50	54 — 81 — 112 — 150 — 264
Noise per 2.54 km s^{−1} channel after imaging		
Noise in residuals [mJy beam ^{−1}]	5.2 — 16 — 22 — 67 — 117	0.8 — 3.7 — 5.5 — 7.1 — 10.6
Peak intensity [Jy beam ^{−1}]	0.11 — 0.42 — 1.4 — 3.2 — 27	0.04 — 0.10 — 0.29 — 0.61 — 1.1
Peak dynamic range	5.4 — 16 — 51 — 116 — 264	7.1 — 21 — 51 — 94 — 189

Notes. These numbers refer to internal release “version 4” constructed with the PHANGS–ALMA pipeline “version 2.0.” This corresponds to the first PHANGS–ALMA public release. We report numbers for the full set of processed data, though some of these are not part of the initial public release because they are archival or still proprietary. See Figures 7–9. Note that some galaxies have been imaged with only the 7 m array, so the samples contributing to the two columns differ. Boldface numbers highlight the 16th percentile, median, and 84th percentile values.

^a Refers to individual mosaics galaxy parts. These are imaged separately and then linearly mosaicked in the image plane (Section 6).

and deconvolution. Section 5 reports our total power processing procedures. Sections 6 and 7 explain our approaches to cube postprocessing and product creation. Section 8 provides an overview of our quality assurance procedures, including end-to-end tests of our pipeline using simulated data. The appendices list the contributions of members of the PHANGS–ALMA data reduction group, and report on tests related to the combination of total power and interferometer data, the stability of the flux calibration in the total power data, and the relative performance of 7 m and combined 12 m+7 m array imaging.

2. Workflow, Definitions, and Implementation

2.1. Workflow

We begin with calibrated $u-v$ data of the sort produced by the ALMA (or VLA) interferometric calibration pipelines. This is stored in the CASA data format of a “measurement set” in which the visibilities have a calibrated phase and amplitude scale. Starting with these data, the pipeline carries out the following steps, which we summarize in Figure 1:

1. *Stage the $u-v$ data.* The pipeline begins by processing the calibrated $u-v$ data into a form appropriate for imaging. It extracts the $u-v$ data associated with the science target and relevant spectral windows from the original measurement sets. Next, it subtracts the continuum signal from the $u-v$ data. It then regrid and rebins all continuum-subtracted, line $u-v$ data onto a common velocity grid to be used in imaging. It also extracts the line-free regions of the spectrum from the original measurement set in order to make a continuum-only $u-v$ data set. This is described in Section 3.
2. *Image and deconvolve the data.* This involves repeated calls to CASA’s `tclean` task interleaved with the creation of masks that guide the deconvolution and checks for convergence. We use a mixture of multiscale

and single-scale CLEAN calls during this process. This is described in Section 4. In parallel, we reduce total power data via the calibration and imaging pipeline presented by Herrera et al. (2020). We summarize these steps in Section 5. There, we also describe the issue of telluric ozone contamination. This issue specifically affects the PHANGS–ALMA CO(2–1) data.

3. *Postprocess the imaged data.* The pipeline applies primary beam corrections, convolves the data to have a round synthesized beam, combines the interferometric and total power data, mosaicks together multipart fields, converts the data to have units of Kelvin, and trims and downsamples the cubes in order to save disk space. Finally, the images are exported into science-ready FITS cubes. These steps are described in Section 6.
4. *Derive additional high-level data products.* The pipeline creates versions of these cubes at several angular and physical resolutions. For each cube and resolution, it creates a noise model that accounts for spectral and spatial variations.

The pipeline uses this noise model to create masks that identify the location of likely signal. We create two sets of masks. The “broad” masks have high completeness, meaning that they include most of the emission in the cube. The “strict” masks have low false-positive rates, meaning that they include only regions where emission is detected at high confidence.

Using these masks, the pipeline produces maps of velocity, integrated intensity, and a suite of other “moments” of the intensity distribution, along with associated uncertainties. This is described in Section 7.

2.2. Definitions

For the most part, this paper uses general radio astronomy terminology and jargon associated with the standard ALMA

Figure 1. Overall pipeline workflow. Schematic view of the pipeline steps. We begin by staging the calibrated visibility data in a form appropriate for imaging (Section 3). We then image and deconvolve the data (Section 4). In parallel, we reduce and image the total power data (Section 5). Next, we postprocess the imaged products into science-ready data cubes and images (Section 6). Finally, we process the images into more advanced science-ready products (Section 7).

data reduction package, CASA. We also define a few pipeline-specific terms here:

1. The pipeline considers “*targets*” to be regions of the sky that will be imaged or processed together. For PHANGS–ALMA, targets are either whole galaxies or parts of galaxies, and each target is a mosaic with tens to more than one hundred individual fields. Within the pipeline infrastructure, each target has an associated mean velocity, velocity width, and phase center.

Some targets only correspond to part of a galaxy. For example, PHANGS–ALMA observed the nearby galaxy NGC 2903 using three separate ~ 150 field mosaics. As described above, this was required due to ALMA’s 150 field limit on any individual observation. We imaged the three observations separately as targets named “NGC2903_1,” “NGC2903_2,” and “NGC2903_3.” These galaxy parts account for the difference between the number of targets and smaller number of galaxies in Table 1.

2. The pipeline makes images for a variety of spectral “*products*.” These are either line products or continuum products. Line products are defined by a spectral line, which sets the rest frequency to be used, and a velocity channel width. For example, CO(2–1) at $\sim 2.54 \text{ km s}^{-1}$ channel width defines the main PHANGS–ALMA line product. C¹⁸O(2–1) at 6.0 km s^{-1} channel width defines another. Continuum products represent the integrated continuum intensity after excluding all user-defined spectral lines of interest in the window.
3. Each input data set is tagged with an “*array combination*.” This does not need to refer to a rigorous antenna or array setup (e.g., ALMA’s C43-1 configuration). The

purpose of the array combination tag is to group data that will be imaged together. For example, PHANGS–ALMA processes data for all main array compact configurations as a single array combination, which we call “12 m.” We also process the ACA 7 m data together as part of an array combination called “7 m.” Finally, we process the ACA and main array data together in an array combination called “12 m+7 m.”

We also define “feathered array combinations.” These combine an interferometric array combination and total power data (“tp”). For PHANGS–ALMA, these are “12 m+7 m+tp” and “7 m+tp.”

2.3. Implementation

As of the PHANGS pipeline “version 2.0” described in this paper, the pipeline consists of a series of linked programs designed to run in CASA (McMullin et al. 2007) and a Python environment equipped with numpy (Oliphant 2006), scipy (Virtanen et al. 2020), astropy (Astropy Collaboration et al. 2013, 2018), and several affiliated packages, most notably reproject,⁴² spectral-cube,⁴³ and radio-beam. Currently, the total power reduction scripts and quality assurance scripts still exist as separate packages.

Both the total power pipeline and our “version 2.0” processing pipeline are publicly available on GitHub.⁴⁴ Our intention is that development will continue on this public version as long as the software remains useful, with “version

⁴² <https://reproject.readthedocs.io/>

⁴³ <https://spectral-cube.readthedocs.io/>

⁴⁴ https://github.com/akleroy/phangs_imaging_scripts and https://github.com/PhangsTeam/TP_ALMA_data_reduction.

2.0” benchmarked as a release. Many of the quality assurance procedures are written in IDL and Python, and are specific to PHANGS–ALMA, so these are not part of the general publicly available pipeline.

We use several different versions of CASA for processing. We note which versions we use for each application in the relevant section. We did not impose a strict version requirement on the *astropy* packages, but mostly used version 4.0 of *astropy*, version 0.4 and after for *spectral-cube*, and version 0.7 and later for *reproject*. We draw the frequencies of spectral lines from *splatalogue* (Remijan et al. 2007).

During prototyping and quality assurance, we also made extensive use of IDL, including the astronomy user’s library (Landsman et al. 1993), *cprops* (Rosolowsky & Leroy 2006), and an updated version called *cpropstoo* (Leroy et al. 2015). For the total power data, we made heavy use of the GILDAS package,⁴⁵ especially CLASS and ASTRO, to prototype, investigate the telluric contamination, and deal with challenging processing cases.

In practice, the PHANGS pipeline is built around a set of modules that are wrapped and called by a series of “handler” classes. The modules contain routines that can run on any input file. They implement tasks like linear mosaicking, spectral line extraction, mask creation, etc. These tasks do not require the rest of the pipeline infrastructure to run, and could be used in other applications. The handlers are aware of the larger project. They interface with user-provided data files, manage directories and files, and loop over targets, spectral line setups, and array configurations. The handlers construct a series of calls to the task-oriented modules in order to implement the steps described in this paper.

The user establishes the input parameters for a project through a series of input text files, which are read and used by the handlers. In these files, the user lists the input calibrated measurement sets and associates each with a target name and array configuration. They also define the targets, specifying a phase center and velocity range for imaging, associating targets that should be linearly mosaicked, and inputting adopted distances to each target. The user inputs also specify the spectral grid, target line, and array combinations for imaging and postprocessing. Finally, the user defines which data products to create, including choosing the angular and spatial scales to be analyzed. In principle, many of these choices could be automated, but we found that leaving them as input parameters worked well for a survey the size of PHANGS–ALMA. In practice, the PHANGS–ALMA choices serve as widely applicable defaults, and most of the customization to define new projects involves simply defining targets, listing input data, and choosing the relevant observed lines.

The pipeline is then executed through a master Python script, either through the shell or a command line call. Staging, imaging, and postprocessing are run inside of CASA. Derivation of data products is run outside of CASA in a pure Python environment. For many applications, the pipeline is trivial to parallelize by simply starting multiple runs targeting different galaxies.

More details and examples can be found with the software itself. The rest of this paper focuses on the procedures used to process the data rather than on the details of the software.

3. Staging of Visibility Data

For each PHANGS–ALMA observation, we apply the observatory-provided calibration and flagging in order to produce a calibrated measurement set. Then, for each combination of target, spectral product, and interferometric array combination, we construct a “staged” visibility data set that will be used in imaging (Section 4). This staged data set combines all relevant visibility data, including data from different ALMA projects, into a single file on a common velocity grid.

The PHANGS–ALMA pipeline assumes calibrated input $u-v$ data. To verify that the input $u-v$ were correctly calibrated, we carried out a by-hand inspection of the calibrated Large Program data. We describe this briefly before discussing the other data processing steps.

3.1. Starting Point

We begin by applying the calibration and flagging produced by the ALMA observatory interferometric pipeline (L. Davis et al. 2021, in preparation) to the data. This step uses the same version of CASA as the original ALMA observatory pipeline run in order to avoid any potential issues arising from changes in calibration tables with CASA version. The ALMA observatory pipeline version changed over the course of the project. Data from the PHANGS–ALMA pilot projects (from Cycles 2 and 3) were mostly calibrated using the Cycle 3 pipeline available with CASA version 4.5.3. Most data from the PHANGS–ALMA large program were calibrated using the Cycle 5 version of the pipeline available with CASA version 5.1.1. Most of the extension projects were calibrated using the Cycle 6 version of the pipeline delivered with CASA 5.4.0.

For PHANGS–ALMA, the ALMA interferometric calibration pipeline performance and observatory quality assurance was excellent. We did not find additional flagging to be necessary, which largely reflects that the data have already been quality assured by the observatory before delivery. To verify this, at several stages during the project, we carried out the inspection described in the next section. These checks aimed to determine whether the pipeline either missed significant flagging or appeared to flag real signal. We did not find any problems serious enough to appreciably affect the final images, so we proceeded using the observatory-provided calibration.

This paper focuses on ALMA observations, but the pipeline also works for other types of data. When we use the pipeline for data with less stringent quality assurance or less stable calibration, the process tends to be iterative. For example, we first image the data. This initial imaging often reveals defects or issues indicating bad data or imperfect calibration. We then improve the flagging, recalibrate, and reimage the data. These flagging and recalibration steps occur outside the PHANGS–ALMA pipeline. After improving the visibility data, the PHANGS–ALMA pipeline is rerun to stage and image the data again. This workflow is common for, e.g., VLA 21 cm data in which radio frequency interference (RFI) can play a large role.

(No) Self-calibration: We did not apply self-calibration to the PHANGS–ALMA data, and we have not yet implemented self-calibration in the PHANGS–ALMA pipeline. The PHANGS–ALMA CO(2–1) images do not appear dynamic range limited, and our mosaic observing strategy does not lend itself to self-calibration. Most fields in most of our sources do not contain bright enough emission to allow for self-calibration. When bright sources are present, they tend to be confined to a small part of

⁴⁵ <http://www.iram.fr/IRAMFR/GILDAS>. For more information about the GILDAS software, see Pety et al. (2005).

Figure 2. Example of the variation of rms noise per pointing in the $u-v$ data. This figure illustrates one of the checks that we carried out during manual quality assurance of the PHANGS–ALMA $u-v$ data (Section 3.2). For each mosaic pointing of NGC 4303, we plot the rms noise in the $u-v$ data as a function of the phase center of that pointing. The left panel shows results for the 7 m observations; the right panel shows those for the 12 m observations. The color scale indicates the variation of rms noise in that pointing from the median value across the map. Maximum deviations are $\sim 3\%$ for the 7 m data and $\sim 12\%$ for the 12 m data. In general, noise correlates with elevation in our data. The mild gradient in decl. in this example reflects the fact that, by chance, for this particular galaxy, the change in observation elevation correlates with decl.

the mosaic, and so are visited only infrequently as part of the mosaic observations.

3.2. Manual Quality Assessment of PHANGS–ALMA $u-v$ Data

As part of the PHANGS–ALMA data reduction process, we inspected the calibrated $u-v$ data from our pilot programs and the Large Program. This inspection focused on the calibrated data, i.e., the direct output from applying the observatory-provided calibration. We inspected:

1. *Observation setup.* We checked the calibrated measurement sets and delivered weblog to confirm that our observational setup was correct. We verified that the observations contained the correct number of fields, total integration time, number of antennas, pointing position on the sky, $u-v$ coverage, and antenna positions.
2. *Observing conditions.* We verified that the weather conditions and related parameters in the weblog were roughly constant across the observations, and that they matched expectations. We checked the precipitable water vapor (PWV), air pressure, humidity, temperature, and wind speed and direction.

We also inspected the antenna-based T_{sys} measurement versus frequency, and compared these to the PWV of the observation. For PHANGS–ALMA CO(2–1) observations, the typical T_{sys} is ~ 70 K, with the highest T_{sys} of ~ 100 K around the weak atmospheric absorption at 231.3 GHz.

3. *Calibrator inspection.* For the pilot program and the first part of the Large Program, we examined the calibrated visibilities for the bandpass and phase calibrators. In this inspection, we aimed to identify outliers and assess the need for additional flagging in the calibrated measurement sets. We plotted time-averaged amplitude and phase as a function of frequency, frequency-averaged amplitude and phase as a function of time, and time- and frequency-averaged amplitude and phase as a function of $u-v$ distance. When we found deviations from the expected behavior in the plots, we manually investigated the $u-v$ data in order to find the cause of the aberrations. This investigation generated a candidate set of additional flagging commands.

Overall, we found that the observatory-provided calibrations yielded calibrated $u-v$ data with few visible pathologies. As described below, our tests suggested that adding additional flagging had negligible impact on the final images. Reflecting this, after the first part of the Large Program, we shifted our manual quality assurance efforts from the $u-v$ data to the imaged data (Section 8). We did not manually inspect the calibrator data for the last part of the Large Program and follow-up programs.

4. *Inspection of synthesized beam.* As an additional check on the $u-v$ coverage of the data after flagging, we created a map of the synthesized dirty beam at the observed CO(2–1) frequency using the CASA tool `imager`. We compared the size and axis ratio of the synthesized beam to expectations based on the $u-v$ coverage before flagging, to further verify that the flagging did not have any pathological impact on the data.
5. *Quality across mosaic.* Finally, we examined the spatial structure of the noise across each mosaic. In particular, we calculated the rms $u-v$ amplitude noise at each individual mosaic pointing, considering all frequencies in the main CO(2–1) science window at the spectral resolution of $\sim 2.54 \text{ km s}^{-1}$. We used this test to check for missing data, e.g., due to flagging or other problems with individual fields. Figure 2 illustrates this check for the 12 m and 7 m data for one galaxy. In this figure, the field-to-field variations in $u-v$ amplitude noise are 0.3% on average and $\sim 6\%$ at most for the 7 m data, and 1.7% on average and $\sim 9\%$ at most for the 12 m data. These results are typical of the targets that we inspected.

The tests described above would often suggest a modest amount of possible additional manual flagging. To assess the science impact of a final round of human flagging on the delivered data, we manually flagged the data for two cases. We then compared the resulting images to those made using no additional flagging.

We chose one galaxy with bright CO emission and one galaxy with faint CO emission for this experiment. Next, we manually inspected the $u-v$ data as described above, identified an aggressive set of flagging commands, and applied

the flags to the data. Finally, we imaged the data with and without this additional flagging. In these tests, the difference between the original and the manually flagged data cubes is less than 5% in total flux and 5% in rms noise for both the bright and faint targets.

In addition to inspecting the quality of individual $u - v$ data, we searched for consistency in the overall calibration scale across the full data set. We examined images of the interferometric calibrators and the flux scale solved for by the pipeline. When we plot the derived flux of any specific secondary calibrator as a function of time, we find good overall consistency among the 12 m array, the 7 m array, and the ALMA calibrator database. For the total power data, we find relatively stable gains, expressed as the observatory-provided Jansky-per-Kelvin (Jy K^{-1} or Jy-per-K) values, across the whole data set at any given time. We did observe that, for observations taking place on similar dates, there is a $\sim 7\%$ difference in the observatory-provided Jy K^{-1} between data delivered before and after the last quarter of 2018 (see Appendix B). As this is an observatory-derived calibration, we accepted the provided values for “version 4” of the PHANGS–ALMA delivery. However, we note that, based on consultation with the observatory, future releases are likely to see the overall flux of some galaxies decline by 2%–5% (see discussion in Appendix B). Only five galaxies have observations delivered both before and after the date of this transition. In Appendix B, we show that, for all other galaxies, the total power data show excellent internal consistency, with rms variation of about $\pm 3\%$.

The inspection steps described above repeatedly verified that the calibration and flagging delivered by ALMA are science-ready, in good agreement with the observatory goals and our previous experience with the telescope. Given the minimal impact of additional flagging, we decided to adopt the observatory-delivered calibration for the PHANGS–ALMA processing.

For the rest of the project, we trusted our detailed quality assurance on the imaged data (Section 8) to reveal any remaining issues with the data.

3.3. Staging and Continuum Subtraction

The PHANGS–ALMA pipeline begins by extracting the calibrated data for each galaxy part and spectral line using the CASA task `split`. We select only the science target, as specified by either the “scan intents” recorded by ALMA, or we manually select a user-provided field or set of fields.

If the continuum subtraction requires multiple spectral windows, we select all spectral windows. Otherwise, for line products, we select only the spectral windows overlapping the line of interest, given the mean recessional velocity and velocity width of the source.

After extracting the science target and window of interest, we subtract the continuum using the CASA task `uvcontsub`. The pipeline is aware of a set of bright lines and the user-provided systemic velocity and velocity width for each target. The pipeline calculates the spectral footprint of each line in the $u - v$ data, and excludes channels that contain line emission when determining the continuum level. For PHANGS–ALMA, we excluded the regions around CO lines from continuum determination. These are the only bright spectral lines in our setup.

For PHANGS–ALMA, we used a single spectral window (spw) for continuum subtraction, which we carried out for all

targets. We fit a polynomial of order zero (i.e., a constant) and fit the continuum in each individual integration. For CO(2–1), the observations used a spectral window with width $\approx 1200 \text{ km s}^{-1}$, and for other lines like $\text{C}^{18}\text{O}(2-1)$ that were covered, the velocity coverage was even larger. This is much larger than the velocity width of any of our targets, leaving wide bandwidth for continuum subtraction. Given the low fractional bandwidth of a 1200 km s^{-1} spectral window and the low signal-to-noise ratio (S/N) of the continuum near the CO(2–1) line, we found that a zeroth-order polynomial did a good job removing the continuum.

For cases with brighter continuum, the pipeline can fit polynomials of higher order, with the order set by the user. In this case, the user can specify the fit to span multiple spectral windows. This is useful, e.g., for ALMA data at Band 7 or above, and for VLA data at L band, where the continuum is strong and the slope is steep. Using multiple spectral windows is also useful when the spectral line of interest covers the entire window, leaving no free bandwidth to fit the continuum. In this case, the pipeline will extract all relevant spectral windows as part of the `split` call above and run `uvcontsub` on all of them, combining spectral windows for the fit.

Time binning: Optionally, at this stage, we also apply some time binning to the data. This is specified by the user when defining each interferometric configuration (e.g., “7 m,” “12 m,” “12 m extended”). This allows the time binning to be defined in a way that avoids time smearing but compresses the data as much as possible. We did not use this option for PHANGS–ALMA, but this is a common step used in VLA 21 cm data processing or processing of ALMA ACA data, especially at Band 3.

3.4. Spectral Regridding and Rebinning

After continuum subtraction, for each spectral line of interest, we create a line-specific measurement set that combines all data on the common velocity grid to be used for imaging. This operation begins with the continuum-subtracted $u - v$ data.

The output spectral grid adopts the “radio” Doppler shift convention, in which $|\delta v|/c = |\delta \nu|/\nu_0$, and we work mostly in the kinematic local standard of rest (LSRK) frame. The user provides the central velocity and width of the final frequency grid as an input. For PHANGS–ALMA, these were initially estimated from large extragalactic databases like NED and LEDA (Paturel et al. 2003; Makarov et al. 2014). We then refined them after inspecting a first round of imaging. On average, the careful systemic velocity estimates using PHANGS–ALMA CO data in Lang et al. (2020) differ from the radio velocity estimates in LEDA by $\sim \pm 5 \text{ km s}^{-1}$ and from the optical velocity estimates by $\sim \pm 10 \text{ km s}^{-1}$. This is small compared to the overall velocity widths used for PHANGS–ALMA cubes. This width for most cubes is $500\text{--}1000 \text{ km s}^{-1}$, with larger values for more massive, heavily inclined galaxies, and smaller values for face-on and low-mass galaxies.

To place the data on the final frequency grid, we first call the CASA task `mstransform` to place all observations onto a velocity grid with a common starting channel and channel width in the LSRK frame. This step converts from the topocentric frame, and so adjusts for changes in the Earth’s motion compared to the LSRK frame. This operation reduces the data to only a moderate velocity range of interest around the line of interest.

After this, we call the CASA task `mstransform` again to rebin the data to the final channel width of $\sim 2.54 \text{ km s}^{-1}$ for PHANGS-ALMA. This rebinning averages together an integer number of channels, typically 5–6 for PHANGS-ALMA CO(2–1) data, and uses no interpolation. The rebinning factor is picked to ensure that the final channel width is as close as possible to the desired spectral resolution for that configuration without exceeding the specified value.

Next, we combine all regridded and rebinned spectral windows for each target and spectral product into a single measurement set using CASA’s `concat` task.

We adopt this regrid-then-rebin approach in order to work around current limitations in CASA’s spectral regridding capabilities, which we describe below. For the PHANGS-ALMA CO(2–1) data, this procedure yields a final channel width, and so a final spectral resolution, near $\Delta\nu \approx 2.54 \text{ km s}^{-1}$ for CO(2–1) and $^{13}\text{CO}(2-1)$, with minor variations from target to target. It yields a resolution near $\Delta\nu \approx 6.0 \text{ km s}^{-1}$ for C $^{18}\text{O}(2-1)$; see Table 1.

After combining the data, the user has the option of reweighting the visibilities by the measured noise using CASA’s `statwt`. This reweighting ensures self-consistent weights in each final line data set, but risks introducing pathologies if real line or continuum emission contaminates the weight calculation. For PHANGS-ALMA, this step occurs after continuum subtraction, so the main danger is contamination by broad line emission. We do apply this procedure to PHANGS-ALMA. We used a 50 km s^{-1} wide window at each edge of the final spectral window for reweighting with `statwt`. This process excludes channels associated with the line itself from the weight calculation. The new weights reflect noise measured from channels far from the systemic velocity of the galaxy.

Noise and spectral regridding in CASA: Our rebinning and regridding strategy introduces some frequency dependence into the noise in the final data products, and also leads to some channel-to-channel correlation. While this is unfortunate, our strategy appears to reflect the best current option, given the spectral regridding capabilities of CASA. We expect that this situation will improve in future versions of CASA. Since it leaves an imprint on our data and likely affects a significant amount of already published ALMA data, we explain the effect here.

The noise pattern arises from the interpolation carried out by CASA’s `mstransform` task. The `mstransform` task can only regrid to channel widths larger than those in the input data. In the case where the output channel width is not an integer multiple of the input channel width, this regridding leads to a varying number of independent data points contributing to different output channels.

We illustrate this effect in Figure 3. We begin with a pure noise, 100 channel visibility data set created using CASA’s `simalma`. The nominal frequency and channel width are $\sim 230 \text{ GHz}$ and 1 MHz , but do not matter for this exercise. In the top panel, we plot the noise spectrum in the original visibility data, which is nearly flat. In the rest of the panels, we show the noise spectrum after regridding the data to new channel widths using `mstransform`.

Figure 3 shows periodic noise variations in the regridded data. The periodicity is set by the fractional difference between the output channel width and an integer multiple of the input channel width. For example, consider regridding to a new grid

Figure 3. Noise and regridding in CASA. An illustration of regridding effects on the noise spectrum of $u-v$ data. We begin with a pure noise data set created by CASA’s `simalma`. This data set produces the noise spectrum seen in the top panel. Subsequent panels show the noise spectrum in the visibility data after regridding using CASA’s `mstransform` task and using `visstat` to measure the noise spectrum. Regridding introduces patterns into the noise, due to the uneven amount of independent data contributing to each output channel. The frequency of the variations is set by the aliasing between the new and original channel width. For example, increasing the channel size by a factor of 1.05 creates a pattern with 20 channel periodicity. To see this, compare the second ($1.2\times$ increase in channel width), fourth ($1.9\times$), and bottom ($1.05\times$) panels. In the bottom right of each panel, we quote the channel-to-channel correlation of the data. This is the linear correlation coefficient between noise values in adjacent channels of individual spectra, and would ideally remain zero under interpolation. The correlation is a function of the statistical independence of the data contributing to the new channel and the interpolation method used. Because cubic interpolation draws on many nearby channels, the correlation induced between immediate neighbors is weaker.

with a channel width 1.2 times the original channel width. During regridding, sometimes a single input channel dominates the data in an output channel. In these cases, other channels do contribute but might only receive, e.g., 20% of the weight in the interpolation. Other times, two input channels are equally weighted and averaged together to form the new output channel. This latter case effectively averages together twice as many independent data points, and will thus have $\sqrt{2}$ times lower noise.

When the output channel is only slightly different from an integer, rebinning the position (in frequency) of output channels relative to the position of input channels “slides”

across the output data set. As a result, the amount of independent data contributing to an output channel varies smoothly across the output data set. The periodicity of the variation is set by the fractional difference between channel size and integer rebinning. For example, when gridding to channels a factor of 1.05 or 1.95 larger than the original channel, the output grid steps are offset by 0.05 initial channels at each channel, and periodicity over 20 channels is expected.

In more extreme cases, the interpolation creates rapid variations and a “sawtooth” pattern in the output noise spectrum. For example, consider gridding from a 1 km s^{-1} channel to a 1.2 km s^{-1} channel. Every ~ 5 output channels, the balance of independent input data will shift from 5-to-1 to 1-to-1 and then back.

In addition to noise variation, the interpolation also affects the correlation between the intensity in successive channels. Because of the variable amount of input data per output channel, the interpolation both introduces channel-to-channel correlation and leads to variations in this channel-to-channel correlation. In optical terms, this processing broadens the line spread function of the data and leads to some dependence of the line spread function on frequency. Figure 3 notes the magnitude of the induced channel-to-channel correlation for each case.

These issues reflect current limitations of CASA, and we expect that the situation will improve in the future. The issue could be addressed by using the `fftshift` option in `mstransform` but that option was not functioning as intended in the versions of CASA that we used. Alternatively, the effect could be mitigated by allowing `mstransform` to oversample the line spread function (i.e., to move to smaller channel width). In this case, heavily oversampled data could be convolved with an appropriate kernel, to produce an even amount of independent data per final, coarser output channel. This functionality is also currently not available.

Regridding in the pipeline: To minimize the effects of the interpolation scheme, the pipeline picks an output channel size that leads to only slow noise variations, i.e., a much more “stretched out” version of the last panel in Figure 3. During the initial regridding step, we increase the common channel size by a small factor, $\epsilon \approx 3 \times 10^{-4}$, compared to the largest channel in any input data set. This will lead to slow noise variations on scales of $\gtrsim 1000$ channels. After this regridding, we rebin the data.

The magnitude of this effect is damped out by the rebinning that follows our regridding. At this stage, many independent channels are averaged together to form each final output channel, e.g., for PHANGS–ALMA CO(2–1), we rebin by a factor of 5–6. As a result of this rebinning, the fractional difference in the amount of independent data in a final output channel varies only modestly across our data.

Still, this effect is enough to induce gradual noise variations with magnitude of $\sim 10\%$ and corresponding variations in the channel-to-channel covariance. These algorithm-induced variations combine with real receiver temperature variations and atmospheric effects to yield the final frequency dependence of the noise in our data cubes (see Section 7).

3.5. Continuum Extraction

We also extract a line-free continuum measurement set. We begin by making a continuum measurement set that includes all spectral windows in each input measurement set. We then cycle

through a list of bright extragalactic emission lines. We use the user-supplied systemic velocity and width to calculate the frequency footprint of each bright line. Whenever a line would overlap the data, we flag all channels associated with the line.

The user can choose which bright lines to consider for flagging. For most PHANGS–ALMA data, we flag only the CO(2–1) and $\text{C}^{18}\text{O}(2-1)$ lines. These represent the only bright lines in our bandpass. This flagging amounts to a flagged bandwidth of $\sim 0.75 \text{ GHz}$ out of the total 6.75 GHz bandwidth observed. For observations in later cycles that cover $^{13}\text{CO}(2-1)$, we also flag that line.

After this flagging, we combine all of these line-free measurement sets using the CASA task `concat`. At this stage, as for spectral line imaging, the user has the option to reweight the combined data set according to the measured rms noise in the visibility data using the CASA `statwt`. This option runs the risk of downweighting regions with bright emission. For PHANGS–ALMA, the S/N in the continuum is extremely low, and the scatter in amplitude for individual $u-v$ data will be determined mainly by the noise in the data. Therefore, we did apply this option for PHANGS–ALMA.

Finally, we use the CASA task `split` to collapse each spectral window in this continuum-only data set to have only a single channel. This step dramatically reduces the overall volume of the measurement set. For PHANGS–ALMA, even after this averaging, the fractional bandwidth of the widest continuum channels is modest, $\lesssim 1\%$, and bandwidth smearing is not a large concern, given the low S/N of the continuum.

3.6. Staged $u-v$ Data

After the staging steps, we have a single, combined visibility measurement set for each combination of target, spectral product, and array combination. These measurement sets are usually significantly reduced in data volume from the input products. For example, for NGC 4303, the calibrated 12 m and 7 m data total $\sim 48 \text{ GB}$, while the staged visibility data set totals 2.4 GB .⁴⁶ They are on the desired spectral grid with appropriate weighting for imaging and deconvolution.

4. Imaging and Deconvolution of Interferometric Data

We use CASA’s `tclean` task to image the calibrated measurement sets and to deconvolve the emission into a “clean” cube or image.

We adopt a two-stage approach to deconvolution that appears well-suited to complex line emission data. First, we run a multiscale deconvolution with a high threshold, corresponding to an S/N of 4, and little or no constraint on where the deconvolution can place components, i.e., little or no “clean masking.” Next, we construct a new, more restrictive clean mask based on the signal in the current cleaned cube. Applying this clean mask, we shift to a standard single-scale deconvolution approach and clean down to a lower threshold, corresponding to an S/N of 1. This deep cleaning ensures a good deconvolution of the numerous small angular scale sources seen in our observations of nearby galaxies.

Throughout this process, we force frequent major cycles. During a major cycle, the model is projected into $u-v$ space

⁴⁶ Presently, the scaling down for 7 m only data is less dramatic because CASA measurement sets include a large pointing table that cannot be removed. This table represents the majority of the data for 7 m observations, but not for 12 m observations.

and subtracted from the visibility data. The residual $u-v$ data are then imaged and used in the next deconvolution step. By comparing the model and data in the $u-v$ plane, we minimize the impact of CASA’s assumption that the synthesized beam, i.e., the interferometric response to a point source, does not vary as a function of position on the sky. In actuality, the synthesized beam can vary across a large mosaic. This variation can introduce minor inconsistencies during the image-plane deconvolution performed in the minor cycles. By frequently projecting back into $u-v$ space, these inconsistencies are mostly corrected. More generally, the major cycle represents a direct comparison between data and model. Frequent major cycles also improve the accuracy of the deconvolution and can help overcome, e.g., limited sampling of the $u-v$ plane due to the lack of significant rotation synthesis in a single ~ 1 hr ALMA observing block.

Periodically stopping and restarting the clean procedure also allows us to check convergence of the deconvolution. We stop the procedure when the fractional change in the model flux with each new clean call drops below 1%. In PHANGS–ALMA, this condition always coincides with the peak residual inside the clean mask approaching the specified threshold: either four times the rms noise for the multiscale case, or one times the rms noise for the single-scale case. When setting these thresholds, the noise level is estimated from the data based on the median absolute value of the residual image. The noise estimated in this way changes relatively little during the course of deconvolution.

This procedure has proven robust. It runs with minimal human intervention across all of PHANGS–ALMA and many other line emission maps. It also works well with many VLA 21 cm HI data sets, though we note a few caveats below. In our view, the key choices were:

1. Use multiscale clean with no clean mask or a very nonrestrictive clean mask and a relatively high S/N threshold.
2. Force many major cycles.
3. Clean deep with a carefully directed single-scale clean, adopting a low S/N threshold.
4. Direct this single-scale clean by applying automated masking to the current deconvolved image, rather than, e.g., the residuals.

The pipeline allows user-input clean masks, but these are not necessary for good performance. When we use clean masks at the multiscale stage, they must be very broad in order to avoid divergence due to interactions between the clean algorithm and the mask boundary. Any user-supplied mask is then used as a prior during the automated creation of the single-scale mask. At this stage, the user-supplied masks help avoid cleaning noise spikes in the often large, signal-free regions of the cube. Avoiding these noise spikes will have a mild impact on the final noise properties, but the main gain is to save computing time during the single-scale clean. Thus, while we do use input clean masks for PHANGS–ALMA, these masks are not crucial to the overall performance of the pipeline deconvolution. Indeed, our first-pass imaging for the PHANGS–ALMA targets without any clean masks yielded almost the same results as the final imaging run. By contrast, supplying an overly restrictive mask often biases the deconvolution and can lead to divergence during multiscale cleaning.

We illustrate the procedure for one galaxy in Figures 4 and 5. Figure 4 shows the deconvolution of the 7 m data for that galaxy. Figure 5 shows the combined deconvolution of the 12 m and 7 m data. Both figures show snapshots of a 20 channel “slab,” i.e., an integral across 20 velocity channels, in one PHANGS–ALMA galaxy. Because the integral extends across the slab, the S/N of these images is improved by a factor of ~ 5 compared to the individual channel maps themselves. Thus, these visualizations show a very aggressive stretch that could bring out artifacts not necessarily visible in individual channel maps.

CASA version: `tclean` refers to the latest CLEAN algorithm implementation available in CASA. This task evolved significantly over the course of the PHANGS–ALMA project. For the “v2 pipeline” and v4 PHANGS–ALMA data release associated with this paper, we imaged the data using `tclean` in its serial (i.e., nonparallel) mode in CASA version 5.4.0.

PHANGS–ALMA CO(2–1) imaging summary: Table 2 and Figures 7–9 summarize our application of this procedure to image the PHANGS–ALMA CO(2–1) data. They report the minimum, maximum, median, and 16th–84th percentile range of key quantities, including the properties of the synthesized beam, the area imaged, and the noise and dynamic range achieved in the cubes. For PHANGS–ALMA, we imaged both the 7 m array data and the combined 12 m+7 m array data. We report numbers for both array combinations, though we emphasize that, when both arrays are available, we strongly prefer the combined 12 m+7 m result to that from the 7 m alone (see Appendix C).

4.1. Imaging

Most of the inputs to `tclean` are tunable parameters in the pipeline. By default, PHANGS–ALMA uses the following imaging parameters:

1. *Cell size.* We use the ALMA observatory-developed `analysisutils` package to estimate the size of the synthesized beam based on the $u-v$ coverage of the data. The pipeline then picks a cell size that is both a round number, e.g., $0''.05$ or $0''.2$, and oversamples the synthesized beam by a factor of $\gtrsim 4$ along the minor axis and more along the major axis. As shown in Table 2 and Figure 7, for PHANGS–ALMA, we place 4–7 pixels along the beam minor axis and 6–10 pixels across the beam major axis.
2. *Image size.* The pipeline chooses an image size with a linear extent $>20\%$ larger than the field of view of the data themselves. We choose an image size in pixels that matches the recommendations for best performance using CASA’s fast Fourier transform (FFT) algorithm, i.e., that is even and can be factorized to 2, 3, 5, and 7 only. As Table 2 and Figure 7 show, for PHANGS–ALMA this translates to typically 240–384 pixels across the ACA 7 m data cubes and 1152–2304 pixels across the combined 12 m+7 m data cubes. The $>20\%$ buffer to the image size can be seen as white space in Figures 4 and 5.
3. *Frequency grid.* For line cubes, the pipeline adopts the frequency grid set during the $u-v$ data processing described above. For the delivered PHANGS–ALMA CO(2–1) imaging, this translates to $\sim 2.54 \text{ km s}^{-1}$ channel width with minor variations from target to target.

Figure 4. Example of deconvolution for PHANGS-ALMA 7 m array data. Integrated emission from a 20 channel thick slab in the CO(2–1) data cube for one PHANGS-ALMA target. The left column shows the image. The right column shows the residuals. From top to bottom: the dirty image; the image after our first stage of cleaning, which uses multiscale clean; and the image after both rounds of cleaning, which follows the multiscale clean with a single-scale clean. Contours show the clean mask at each stage. All images share the same stretch, which saturates at 10 times the rms noise level in order to focus the image on the low-level, extended structure rather than the bright, easily cleaned sources. Along with Figure 5, this figure illustrates our approach to imaging and deconvolution. White space indicates masking used in imaging. The second row shows the broad, inclusive nature of the user-supplied clean mask used in multiscale clean. The third row shows the much more restrictive mask automatically generated and used in single-scale clean. The middle right panel shows the filamentary, clumpy nature of the residuals after multiscale clean. The bottom right panel shows the clean residual image after single-scale clean.

4. *Gridding algorithm, weighting, and primary beam cutoff.*
By default, the pipeline uses CASA’s “mosaic” gridding algorithm and weights the $u-v$ data according to the

“Briggs” scheme. It defaults to robustness parameter $r = 0.5$, which offers a good compromise between noise and resolution. By default, it images out to a primary

Figure 5. Example of deconvolution for PHANGS-ALMA combined 12 m and 7 m array data. Similar to Figure 4, but now showing results for imaging the 12 m+7 m array data together.

beam cutoff of 0.25. We adopt all of these parameters when imaging PHANGS-ALMA.

Following the observatory recommendations, we set `mosweight` to True and calculate the $u-v$ weighting for each field separately. Following the documentation, this can improve imaging performance for mosaics at the expense of a slightly larger beam. Because we imaged in CASA 5.4, before the `perchanweightdensity` parameter was introduced, our

imaging effectively sets `perchanweightdensity` to False. This parameter instructs `tclean` to weight each channel individually. Similarly to `mosweight`, it should lead to better imaging performance at the expense of a slightly larger beam size. In future runs of the PHANGS-ALMA imaging pipeline using CASA version 5.5, the user can choose whether to adopt per-channel weighting by setting the `perchanweightdensity` in the `clean` call.

Figure 6. PHANGS-ALMA clean mask in projection. A two-dimensional projection of one PHANGS-ALMA “user supplied” clean mask is shown. The background image shows peak intensity along the R.A. axis in the 12 m+7 m imaging for NGC 4303. The black-and-white contours show the locations along the line of sight where the mask is “True” for at least one pixel along the R.A. axis. The clean masks are created based on previous rounds of imaging. This figure illustrates how the clean masks broadly circle emission in the cube, rather than applying any substantial restriction. In order to avoid any edge effects, this mask even reaches slightly beyond the ALMA coverage, hence the extension into the white region. Their main effect is to save computing time by avoiding processing the signal-free regions of the cube. The central rectangle shows a region extending over the full velocity width of the cube centered on the galaxy center. We include a feature like this for all galaxies with bright centers.

5. *Independently image mosaics observed separately.* For PHANGS-ALMA, we observed some galaxies in multiple parts. Each part corresponds to a ~ 150 field mosaic, and the parts were observed separately. We imaged each separate part independently. The choice to independently image each separately observed mosaic is important for PHANGS-ALMA. When we observed a galaxy using several adjacent mosaics, these mosaics were sometimes observed at different times and even different array configurations. This implies a spatially variable synthesized beam across the field, and CASA cannot currently account for position-dependent synthesized beams. Our initial attempts to jointly image multiple large mosaics frequently resulted in divergence. This problem was resolved when we shifted our strategy to image each part separately and then linearly mosaic the parts together.
6. *Dirty image and clean mask alignment:* As the first step in imaging, we constructed a “dirty” cube. This cube used our adopted imaging parameters, but we performed no deconvolution.

If the user supplied a clean mask, as was the case for PHANGS-ALMA, then at this stage we used CASA’s `importfits` and `imregrid` tasks to align the clean mask to the astrometric grid and axis order of the dirty cube.

Slabs, i.e., integrals over 20 successive channels, in PHANGS-ALMA dirty cubes appear as the top row in Figures 4 and 5. As expected, these dirty images look highly distorted due to spatial

filtering through the incomplete $u-v$ coverage of the interferometer. The imprint of the user-supplied clean mask for PHANGS-ALMA appears as a contour in the second row.

4.2. Deconvolution

The pipeline uses `tclean` to deconvolve emission and create a clean cube or image. As described above, this has two main stages: a “wide” multiscale clean and a “directed, deep” single-scale clean. We follow a few general principles in both stages:

1. *Force frequent major cycles.* The pipeline requires “major cycles” to happen frequently. During a major cycle, the approximate image-plane deconvolution is projected back into visibility (Fourier) space and the model is properly subtracted from the data. While computationally expensive, this process produces a more correct residual image, allowing for a more stable, precise deconvolution.

In practice, the pipeline enforces major cycles within each `tclean` call in two ways. First, it limits the number of “iterations” allowed before forcing a major cycle using the `cycleniter` keyword. Second, it uses a combination of `cyclefactor` and `minpsffraction` to set an aggressive threshold for triggering a major cycle. Once the data are cleaned so that the maximum residual approaches this threshold level, `tclean` triggers a major cycle. For PHANGS-ALMA, our default values for these parameters were `cyclefactor` = 3.0 and `minpsffraction` = 0.5. These imply that the threshold is never lower than 0.5 times the peak residual or three times the maximum sidelobe level times the peak residual. By default, the pipeline also uses `maxpsffraction` = 0.8 in order to ensure that some emission is deconvolved in each cycle.

2. *Multiple `tclean` calls with more components deconvolved in later calls.* The deconvolution involved many repeated calls to `tclean`. When the pipeline initially calls `tclean`, it allows only for a small number of clean components, with the number set via the `niter` keyword. It also allows for only a limited number of components to be cleaned per channel before enforcing a major cycle. This is set via the `cycleniter` keyword. Once the overall number of allocated clean components is exceeded, `tclean` stops. Stopping and resuming `tclean` forces a major cycle.

Over the course of the first five `tclean` calls, the pipeline increases `niter` and `cycleniter`. By default, the pipeline increases `niter` by a factor of two each step. It linearly increases `cycleniter`, starting at 100 and increasing it by 100 at each step in the loop. The choice to limit the number of components in any individual call to `tclean` is part of our strategy to trigger frequent major cycles.

This gradual increase in allocated clean components resembles the approach used to create the PdBI CO image of M51 by Pety et al. (2013). The numerical choice of how to progressively increase the number of iterations is *ad hoc*.

3. *Check for convergence between clean calls.* These repeated calls to `tclean` allow us to check for convergence in the deconvolution. After each call and before the next one, we calculate the sum of flux in the model (i.e., the clean components). We compare this flux to the previous model flux to calculate the fractional change in flux and the gain in flux per allocated clean

Figure 7. *Imaging properties related to the beam.* Properties of the imaged CO(2–1) PHANGS–ALMA cubes for the ACA 7 m data only (left column) and the combined 12 m and 7 m data (right column). From top to bottom, we show the FWHM major axis of the synthesized beam, the position angle (measured north through east) of the major axis of the synthesized beam, the beam elongation (defined as major over minor axis), and the elongation as a function of decl. of the source. See Table 2. Note that some galaxies have been imaged with only the 7 m array, so the samples contributing to the two columns differ.

component. When the fractional change in the model flux drops below some threshold, usually 1%, we terminate that stage of the deconvolution and move to the next one. In the case of the multiscale clean, we move to automated masking and single-scale cleaning. In the case of single-scale cleaning, we finish the deconvolution and move to postprocessing.

4. *Common restoring beam.* By default, we use a common restoring beam, meaning that `tclean` restores deconvolved emission with a single elliptical, Gaussian beam across all planes of the cube. The alternative offered by CASA is to track the beam per plane, reflecting differences in how the $u-v$ coverage maps to angular

scale as the frequency changes. For PHANGS–ALMA, the fractional bandwidth, $\delta\nu/\nu$, across our cubes is modest, always < 0.0035 . As a result, the synthesized beam does not change much with frequency, and we do not keep track of a beam per plane. This choice can be changed by the user. For example, a change may be required when many data are flagged in a few channels, which would otherwise result in a large common beam.

4.3. Multiscale Clean

In the first stage of deconvolution, we employ the CASA implementation of the “multiscale” deconvolution algorithm

Figure 8. Imaging properties related to the mapping area and image size. Properties of the imaged CO(2–1) PHANGS–ALMA cubes for the ACA 7 m data only (left column) and the combined 12 m and 7 m data (right column). From top to bottom, we show the number of pixels across the FWHM of the beam minor axis, the number of pixels across the major axis of the cube, the area of sky imaged, and the spatial dynamic range of the imaged region. See Table 2. Note that some galaxies have been imaged with only the 7 m array, so the samples contributing to the two columns differ. Also note that individual mosaic parts are imaged separately. The final spatial dynamic range of those images will be higher than shown here.

(Cornwell 2008). For this stage, PHANGS–ALMA uses a broad clean mask supplied by the user, but the operation also works well with no mask. The scales to be cleaned are also specified by the user as part of defining the configurations. We follow the CASA recommendation regarding choice of scales and use scales from the beam size to within a factor of ~ 2 of the largest recoverable scale.

Multiscale clean includes a tuning parameter, `smallscalebias`, that can be used to bias the results toward small or large scales. We set `smallscalebias` to 0.9 by default, indicating a preference for small scales. During development, we experimented with scales from 0.4 to 0.9. We found higher values less likely to yield divergence. Note that these tests used

earlier versions of CASA, mostly 4.5 and 4.7. This may reflect the common presence of a few bright, clumpy structures in our CO maps.

For PHANGS–ALMA, when deconvolving only 7 m data, we employed scales of $0''$ (i.e., a point source), $5''$, and $10''$. When deconvolving the combined 12 m+7 m data, we considered scales of $0''$, $1''$, $2''$, $5''$, and $10''$. When deconvolving only 12 m data, we used scales of $0''$, $1''$, $2''$, $5''$, and $5''$. These deconvolution scales correspond to the size of round Gaussian clean components before convolution with the dirty beam.

We impose a threshold of four times the rms noise on the multiscale cleaning process. For this purpose, we take a single

Figure 9. *Imaging properties related to noise and dynamic range.* The properties of the imaged CO(2–1) PHANGS–ALMA cubes for the ACA 7 m data only (left column) and the combined 12 m and 7 m data (right column). From top to bottom, we show the rms noise in the residuals of each cube, the peak intensity in each cube, and the maximum dynamic range (peak intensity over rms noise) in any channel of the cube. See Table 2. Note that some galaxies have been imaged with only the 7 m array, so the samples contributing to the two columns differ.

robustly estimated noise value to describe the whole cube (but see Section 7.2). When the peak value in the residual map for each channel falls below this level, cleaning stops in that channel. We estimate the noise from the residual cube, and update this noise estimate between calls to `tclean`. Because we use a robust noise estimator and the cubes contain a large amount of empty volume, the estimated value of the noise changes little between calls. We found that adopting lower S/N thresholds for the multiscale clean led to divergence in the deconvolution (for similar conclusions using VLA data, see Koch et al. (2018b)).

As described above, after each call to `tclean`, we sum the total flux in the model image, i.e., the sum of deconvolved flux. When this flux changes by $< 1\%$ between subsequent calls to `tclean`, we move to the next stage of the deconvolution. Usually, this convergence coincides with the peak residual approaching the S/N-based threshold. If the deconvolution has not converged, then we increase the `niter` and `cycleniter`, and we continue the multiscale deconvolution with a new call to `tclean`.

4.4. Masking and Single-scale Clean

After the multiscale deconvolution converged, there were often still significant residuals around the brightest sources. At this stage, we proceed by deconvolving with the classic, single-

scale CLEAN algorithm (Högbom 1974) and use it to clean down to a threshold equivalent to an S/N of 1. We also generate and apply a much more restrictive clean mask at this step. This masking avoids spending large amounts of effort cleaning signal-free regions of the data cube, and makes it possible for the deconvolution to clean very deeply in regions with signal. The shift to the single-scale clean avoids potential pathological interactions between this more restrictive clean mask and large cleaning scales.

We use the resultant multiscale deconvolved image to construct an S/N-based mask. To do this, we estimate a characteristic rms noise in the cube based on the median absolute deviation of the whole residual cube. Then, we create a mask that includes all regions that have $S/N > 4$. We then expand this mask to adjacent regions with $S/N > 2$. Finally, we extend the mask by one channel in each velocity direction. If the user supplied a clean mask, then during this step we only include pixels in the mask that also lie inside the original clean mask.

In this way, we focus the single-scale clean on regions where signal is already evident in the cleaned maps after the multiscale clean. We note that this approach differs from the automated masking within the `tclean` task in CASA “auto-multithresh.” CASA’s algorithm builds a clean mask based on the current residual emission as part of the major cycle

(Kepley et al. 2020), while we construct a clean mask based on the deconvolved emission outside the deconvolution process. Based on experimentation, we found by eye that our approach did a good job of identifying the regions of the residual image where one would want to clean deeper. Put another way, we use the single-scale clean to “dig deeper” in order to ensure a full deconvolution of already visible bright regions.

During this single-scale deconvolution, we impose an S/N threshold of 1, again using a single robustly estimated noise value to describe the whole cube. This threshold means that we stop the deconvolution in each channel when the maximum residual in that channel reaches a value equal to the noise level. This limit is much lower than the threshold that we adopted for the multiscale clean. This change causes the single-scale clean to deconvolve a large network of filamentary S/N $\lesssim 4$ residuals commonly remaining after the shallow multiscale clean.

As with the multiscale deconvolution, during this step we allocate only a limited number of iterations to each `tclean` call. Between calls, we check for convergence. Again, we define this as the flux in the model changing by $<1\%$ between successive clean calls. We begin these convergence checks after three calls to single-scale `tclean`. This delay allows us time to allocate enough iterations to allow some expectation of convergence.

Our peak residual threshold in individual calls to `tclean` interacts with our fractional-change-in-flux criteria. In practice, the fractional change in flux drops below 1% when the peak residuals inside the clean mask approach the threshold. For PHANGS–ALMA, the single-scale clean thus effectively cleans down to a peak S/N = 1 in the residuals within the clean mask.

4.5. Input or Iterative Clean Masks

As discussed above, user-input clean masks are optional in our approach. Indeed, they mostly do not appear necessary. We imaged every PHANGS–ALMA target without a user-supplied mask before imaging them with masks. These initial images generally appeared similar to the final ones.

The procedure works without an input clean mask because the high threshold adopted for the multiscale clean makes heavy cleaning of noise spikes unlikely. After this, the pipeline creates a clean mask and our automated masking procedure generally appears to work well. The main gains in using a mask appear to be related to performance. Our single-masking approach will still produce some false positives when applied to large signal-free regions. When we supplied broad clean masks that restricted clean to the general area of the galaxy, we avoided time cleaning spurious “islands” of emission during both clean stages.

When provided, input clean masks need to encompass all real emission and be extended compared to the scales used by the multiscale deconvolution. In PHANGS–ALMA, our general procedure is to adopt an iterative approach. We image a target without any prior clean mask. We then convolve the initial deconvolved cube to coarser resolution. Next, we adopt a masking approach similar to that used in product creation below. Finally, we dilate the mask by several channels in the velocity dimension and by about the largest recoverable scale in the spatial dimension.

Specifically, we created our clean masks by convolving the initial 7 m imaging to coarser angular and spectral resolution, $33'' \times 20 \text{ km s}^{-1}$. We constructed a mask at this low resolution

via sigma-clipping. For any galaxy deemed to have a bright central region, we extended the mask over the inner $40''$ diameter to cover the full velocity range of the cube. We found that this was necessary to ensure complete coverage of any compact, high-velocity material associated with the inner disk or outflows. We inspect each mask on a high stretch in all projections of position–position–velocity space, to ensure that the mask includes all emission with enough room for the multiscale deconvolution to place large components.

4.6. Comments on PHANGS–ALMA Imaging

Table 2 and Figures 7–9 summarize the application of these algorithms to the PHANGS–ALMA CO(2–1) data.

Imaging only the ACA 7 m data yields synthesized beam sizes mostly in the range of $6''.8$ – $7''.9$. The beams for the ACA 7 m data tend to be significantly elongated, with the major-to-minor axis ratio typically in the range of 1.4–2.0. The elongation is mostly along the east–west direction and worst at intermediate decl., as expected based on the information provided in the ALMA Technical Handbook, which reports large beam elongations for $-40^\circ < \text{decl.} < 0^\circ$.

For the 7 m imaging, we typically place seven pixels across the major axis of the beam, 4.4 pixels across the minor axis of the beam, and image a cube 240–384 pixels across. On average, the maps are a few square arcminutes in size, with a median 8.2 arcmin^2 . Across the entire 7 m portion of the survey and including archival data, we mapped about 0.4 square degrees. The typical spatial dynamic range of an individual 7 m image, defined as the number of resolution elements along one dimension of the image, is about 28.

For the 7 m imaging, we achieve a typical rms noise of 22 mJy beam^{-1} per 2.54 km s^{-1} channel. The peak dynamic range, meaning peak intensity in a channel divided by rms scatter in that channel, varies across the sample but is mostly in the range of 16–116. Note that this is the dynamic range in an individual channel. The $\sim 2.54 \text{ km s}^{-1}$ channel width places about 5–10 elements, and sometimes many more, across a typical emission line. As a result, the line-integrated S/N is even higher.

The combined 12 m and 7 m imaging typically yields a beam size of $1''.0$ – $1''.6$, with median of $1''.2$. These beams tend to be less elongated, with a median major-to-minor axis ratio of 1.2 (consistent with the expected beam shape based on the ALMA configurations). As with the 7 m data, the elongation tends to place the major axis in the east–west direction. Here, we place about six pixels along the minor axis of the beam when imaging.

The 12 m+7 m cubes are much larger in pixel units, typically 1152–2304 pixels across. Again, the cubes tend to cover a few square arcminutes, usually 2.8 – 7.8 arcmin^2 and 6.2 arcmin^2 on average. The slightly smaller mapped area reflects the larger primary beam of the 7 m antennas and that the 7 m sample includes several very large, nearby galaxies, e.g., NGC 0253, that we did not map with the 12 m array. The total area mapped by the 12 m survey is about 0.15 square degrees, about half the area covered by the 7 m survey.

The spatial dynamic range of the combined images is much higher than for the 7 m only data. The typical spatial dynamic range of 112 corresponds to $>10,000$ independent spectra per image.

The typical noise in the residuals of the combined data is $5.5 \text{ mJy beam}^{-1}$ per 2.54 km s^{-1} channel, and the peak

dynamic range is similar to that in the 7 m only images: ~ 50 on average. Again, the integrated S/N in the maps will be even higher.

We consistently deconvolve more flux when imaging the combined 12 m+7 m array data than when imaging the 7 m array data for the same target. In Appendix C, we analyze this effect using both our full data set and simulated data, in which the correct sky image is known *a priori* (see Section 8.3). Our analysis suggests that this discrepancy is a general feature of ALMA observations of nearby galaxies: compact 12 m array observations play an important role in achieving a complete deconvolution of emission, even when 7 m array observations are present.

4.7. Limitations of the Imaging Approach

Overall, this imaging scheme has proven robust and we have successfully applied it to a variety of ALMA and VLA line and continuum data. However, we have encountered a few cases where the approach needs modification or does not work, and we note these here. First, when imaging sources with bright, not-yet-subtracted continuum emission, our convergence tests need modification. The convergence test focuses on the fractional change in flux. Including one or more high-flux point sources can skew the imaging to converge before any surrounding faint emission has been imaged. More generally, our convergence criteria need to be refined to reflect the desired dynamic range. Our adopted criteria work well for the dynamic range of ~ 10 –1000 expected for PHANGS–ALMA and VLA 21 cm imaging of nearby galaxies.

Second, when imaging structures with extended, highly asymmetric structure, the use of large, symmetric multiscale clean components can lead to oversubtraction. To some degree, `tclean` can make up for this by adding negative components to the model. However, in some cases, either adjusting the `smallscalebias` tuning parameter to emphasize small scales or adopting a more restrictive clean mask can improve performance. We have mainly encountered this issue in applying the algorithm to 21 cm imaging of Local Group galaxies, where extended, asymmetric emission extends across very large scales.

Third, we made several choices in constructing the imaging algorithm. We chose the S/N threshold for the single- and multiscale clean, as well as various gridding parameters, the set of scales for multiscale clean, and details of masking. In principle, the PHANGS–ALMA pipeline can be used to conduct a full regression analysis, exploring the uncertainty associated with changing each parameter within a reasonable range. In practice, because it takes roughly a full day for a server with 24 CPUs and 256 GB of memory to process a typical target, we are only able to carry out a limited number of these tests. In Section 8.3, we describe how we run two targets at multiple S/N levels through complete end-to-end tests of the pipeline. In Appendix D, we carry out a similar test to investigate approaches to short-spacing correction. These tests are already helpful, but due to practical considerations, we have delayed a comprehensive assessment of the uncertainties associated with the choice of imaging parameters to the future.

5. Calibration and Imaging of Total Power Data

We process total power data in parallel with the interferometer data using a separate pipeline. For this, we use the

modified version of the ALMA total power pipeline presented by Herrera et al. (2020). We give an overview of the procedure here, and refer the reader to Herrera et al. (2020) and the publicly available scripts⁴⁷ for more details. We also highlight one specific issue important to the PHANGS–ALMA total power data, namely the contamination of a subset of our data by a telluric ozone line at 229.575 GHz.

This total power pipeline employs a combination of the CASA, GILDAS, and R software packages. Unless otherwise noted, we carry out these steps in CASA version 4.7.2.

5.1. Calibration

We import the single-dish data from the observatory-provided ASDM format to the “measurement set” format. Then, we split the data by antenna and write them into the ASAP data format as “scantables.” Next, we compute and apply the “chopper wheel”-based temperature scale using CASA’s `sdcal2` task. This task calculates the temperature scale from the hot and cold loads plus sky observations (Penzias & Burrus 1973). We use this same task, `sdcal2`, to subtract the “OFF” spectrum from each on-source spectrum. These “OFF” spectra are obtained by integrating on empty sky near the source, so the result is a set of calibrated, sky-subtracted spectra.

5.2. Baseline Fitting

After calibration, we convert the frequency and velocity scales of the spectra from the observatory into the LSRK frame, around the systemic velocity of each galaxy. This step suffers from the same issues regarding CASA regridding described in Section 3.4. These currently represent an unavoidable limitation of the software.

We start by extracting a wide part of each spectrum centered on the systemic velocity of the galaxy. We then fit first-order baselines to the line-free regions of each of these calibrated, sky-subtracted spectra. Baseline offsets and frequency-dependent baseline fluctuations are a common feature of single-dish data. They reflect imperfect matches between the “ON” and “OFF” spectra and instabilities in the receiver, sky, or other parts of the signal path. The fitted baselines will also include any genuine continuum emission from the galaxy. As a result, this step removes any sensitivity to the total power data to continuum.

For simplicity, we define the “line-free” region to be fitted by excluding a fixed velocity range from all spectra in each data set for the baseline fitting procedure. We choose the excluded velocity range to be large so that it easily encompasses all emission from the galaxy. The excluded velocity interval ranges between 200 and 500 km s^{−1}, depending on the target. After fitting, we subtract the fitted baselines from the calibrated spectrum. This procedure is carried out independently for each ALMA execution block and for each antenna.

5.3. Unit Conversion and Combination

After calibration and baseline subtraction, we first apply the antenna efficiency factor provided by the observatory as part of the delivery to convert the intensity scale from units of antenna temperature, in Kelvin, to Jy beam^{−1}.

⁴⁷ https://github.com/PhangsTeam/TP_ALMA_data_reduction

The observatory regularly measures these efficiencies by combining the total power antennas with interferometric observations by the 7 m antennas to provide time-dependent conversion factors. This is done on a per antenna and per observation basis. This ensures a highly reliable flux calibration of the total power data.

In Appendix B, we verify that the individual PHANGS–ALMA total power observations for the same galaxy show only 3% rms scatter in amplitude scale from observation to observation. This is consistent with high-quality overall calibration of the ALMA total power data. We make no additional corrections here, nor do we scale the data during combination with the interferometer data.

Last, we merge the data from all observations and antennas into a single CASA measurement set using the CASA task `concat`.

5.4. Imaging

We grid the calibrated, sky-subtracted, baseline-corrected spectra into a data cube. To do this, we use CASA’s `sdimaging` task, which convolves the irregularly sampled spectra onto a regular grid (e.g., Mangum et al. 2007). For the CO(2–1) data, this convolution uses a spheroidal gridding kernel with a support diameter of 12 pixels and a pixel size of $\sim 2''$.

5.5. Inspection and Quality Assurance

We perform basic inspection by examining the integrated spectra for each scan and each antenna. Occasionally, individual scans or antennas reveal isolated artifacts and are flagged. Less than 1% of our total data were removed in this manner. We also check the “line-free” region from which to calculate the baseline fit, and adjust it if it overlaps with any galaxy emission. If any flagging or baseline region selection changes, we rerun the entire pipeline for the target.

After gridding, we also visually inspect the cubes. Except for the telluric contamination discussed below, they showed no signs of residual pathological spectra or artifacts.

5.6. Telluric Ozone Contamination of CO(2–1) Data

The CO(2–1) total power data for six PHANGS–ALMA targets is contaminated by a spurious line feature of $\sim 50 \text{ km s}^{-1}$ width that peaks near $V_{\text{LSRK}} \approx 1250 \text{ km s}^{-1}$. We ascribe the observed contamination to a relatively weak telluric ozone line at 229.575 GHz rest frequency (i.e., offset by $+1253 \text{ km s}^{-1}$ from the rest frequency of the CO(2–1) line).

In our original total power observations, the OFF position was fixed in the equatorial reference frame, and the contamination affected even more galaxies. In these cases, the feature typically appeared either positive or negative throughout each entire image. Most of the affected targets were then reobserved, this time using an OFF position at the same elevation as the target, i.e., using a fixed offset in azimuth rather than a fixed offset in R.A. and decl.

These reobservations improved the situation, reducing the strength of the feature or even suppressing it entirely. In cases where the feature persisted, the reobservations tended to shift the nature of the contamination. Rather than having a fixed sign across the whole data set, in sets observed with a fixed-elevation OFF, the contamination shifted from positive to negative on opposite sides of the target.

This behavior can be naturally expected from the calibration procedure.⁴⁸ The ON-OFF subtraction used to remove atmospheric emission from the source will leave a remnant contribution proportional to the difference in airmass between the ON and the OFF spectra. To first order, this difference will be proportional to the offset in elevation. This also explains why the interferometric data are not affected by the contamination, beyond a potential mild increase in noise at these frequencies, as they do not reference to a displaced OFF position.

The subsequent baseline subtraction typically uses a first-order polynomial fit. This fit can remove any residual continuum emission, which will vary smoothly and slowly as a function of frequency. However, the baseline fit cannot remove a narrow line feature like the ozone line. The situation becomes even worse when the ozone line overlaps the velocity range covered by the galaxy. Then, the line emission from the sky and the source can become confused.

We tested this scenario by looking for the primary direction along which the ozone feature varied. We found that, as expected, the strength of the feature tended to vary almost linearly along a direction close to the elevation axis at the time of the observations.

We measured the gradient of that linear trend and found that the peak of the ozone feature typically varied by $\sim 0.02\text{--}0.03 \text{ mK arcsec}^{-1}$, with the calculation done in $\sim 10 \text{ km s}^{-1}$ channels. This measured gradient is consistent with an order-of-magnitude estimate using the ATM atmospheric model (Pardo et al. 2001) distributed along with the GILDAS software. The exact value of the gradient appears to depend on elevation and atmospheric conditions. In the most extreme case, it reached four times this typical value.

To the best of our knowledge, contamination by the 229.575 GHz ozone line has not been reported in previous extragalactic, single-dish CO(2–1) surveys, even large mapping surveys covering comparable area to PHANGS–ALMA (e.g., the IRAM 30 m HERACLES survey; see Leroy et al. (2009)). Our best estimate is that this contamination simply reflects the much better sensitivity in the PHANGS–ALMA data compared to previous mapping surveys. The rms noise of our total power maps is typically $2.5\text{--}3.0 \text{ mK}$ per 2.5 km s^{-1} , compared to $\sim 25 \text{ mK}$ per 5.2 km s^{-1} channel in the HERACLES maps.

Results on the telluric contamination have been reported back to the ALMA observatory in a memo by A. Usero et al. This memo recommends observing strategies that can mitigate the effect of the ozone line. In general, the contamination is stronger at lower elevation and the linear trend becomes significantly steeper below $\sim 45^\circ$.

5.7. Strategy for Fitting and Removing Telluric Ozone Contamination

In six galaxies, the telluric ozone line overlaps the CO(2–1) line velocity range and reobservations using a fixed-elevation OFF did not solve the problem. We thus developed a custom procedure to remove the ozone contamination. This procedure, which is currently implemented in the R programming language (R Core Team 2015), works as follows:

1. The procedure operates at the level of individual “execution blocks” (EBs), i.e., individual observing

⁴⁸ For the relevant information for CASA 4.7.2, see Equation (8.1) in Section 8.5.2 of <https://casa.nrao.edu/docs/cookbook/index.html>.

sessions. Because these sessions are relatively short, $\sim 1\text{--}1.5$ hr, we approximate the transformation from the azimuth-elevation frame to the celestial frame as constant and work with the post-gridding cube data for individual EBs. We also assume that atmospheric conditions are stable over this short time, such that the signature of the telluric contamination is constant.

We model the strength of the ozone line at a sky position $\mathbf{x}$ and velocity v as

$$T_{\text{O}_3}(\mathbf{x}, v) = L(\mathbf{x}) \times P(v), \quad (1)$$

where L is a linear gradient that models the amplitude as a function of elevation, and P is the spectral profile of the line. Because we work in the LSRK velocity frame, the peak velocity of the ozone profile P will typically be offset from $+1253 \text{ km s}^{-1}$ by the difference between the topocentric velocity and LSRK rest frame. We calculate the expected offset using the `ASTRO` program of the `GILDAS` software. The velocity difference between the frames varies across EBs by as much as 30 km s^{-1} .

2. For each EB, we generate a *contaminated* CO(2–1) cube with our standard total power pipeline. The only modification is to exclude an additional velocity range around the ozone line during baseline fitting.
3. To determine L in Equation (1), we build a map of the mean intensity in the cube within a $\pm 25 \text{ km s}^{-1}$ range centered on the expected peak velocity of P for the ozone line. We then manually define a two-dimensional mask that encompasses the real CO emission from the galaxy in this velocity range. The signal outside this spatial mask will mostly represent telluric emission/absorption. We fit the unmasked position data as a linear function of R.A. and decl. using a noise-weighted least-squares method to generate an estimate of L .
4. Our model (Equation (1)) assumes that the spectral profile of the line, P , does not vary across the map. To determine $P(v)$ in each velocity channel, v , we first build a mask that encompasses all real CO emission, using a dilated mask technique (similar to the masks discussed in Section 4.4). We then calculate the average (T/L) outside the mask, weighting it by $(L/\sigma)^2$. Here, T and σ are the measured intensity and the rms noise at the corresponding position and velocity channel, respectively. We smooth the initial estimate of P with a five-channel boxcar kernel to reduce uncertainty due to noise. We also set $P(v) = 0$ beyond $\pm 50 \text{ km s}^{-1}$ from the expected peak velocity of the ozone line. This limit ensures the correction does not create any artifacts at velocities where contamination would be negligible even at our high sensitivity.
5. We build a contamination cube from $L \times P$ and subtract it from the original CO cube, to get its *contamination-corrected* version.
6. Finally, we co-add the *contamination-corrected* cubes from all EBs, weighted by their average rms noise, to produce the final total power cube for that galaxy.

As illustrated in Figure 10, this procedure effectively removed signatures of ozone contamination in the six remaining affected PHANGS–ALMA targets. This approach should be useful for sensitive on-the-fly observations of external galaxies with source velocities in the range $\sim 1100\text{--}1400 \text{ km s}^{-1}$. The procedure could also be generalized

to deal with any telluric contamination of on-the-fly mapping maps.

6. Cube Postprocessing

After imaging and deconvolution, we process the interferometric cubes into a final “science-ready” form. This has six main steps: primary beam correction, convolution to a round synthesized beam, linear mosaicking to combine multipart galaxies, feathering to combine interferometric and total power data, downsampling of cubes, and conversion to a Kelvin intensity scale.

6.1. Primary Beam Correction

First, we created a version of each cube that was corrected by the combined primary beam response of all mosaic pointings, B , in each channel. To do this, we use the `CASA` task `impbcor`, which divides the image cube by the combined primary beam response map output by `tclean`. A byproduct of this correction is to increase the noise near the mosaic edges, where B is low.

6.2. Convolution to a Round Beam

The imaging yields elliptical synthesized beams, e.g., as shown in Table 2 and Figure 7 for PHANGS–ALMA. The beam does not tend to align in any useful way with galactic structure, and the elliptical shape only makes analysis more complex. Therefore, as part of postprocessing, we used the `CASA` task `imsmooth` to convolve each cube to a final, round, Gaussian-shaped beam.

In addition to increasing the minor axis of the beam, `imsmooth` required us to slightly pad the major axis of the beam in order to find a viable convolution kernel. This increased the major axis beam size by a small amount, $\lesssim 10\%$. In principle, this resolution loss can be avoided by constructing an appropriate kernel in the Fourier domain. This kernel would have infinite width (in Fourier space) along the major axis in order to avoid convolution in that direction. `CASA` currently lacks this capability, so we instead pad the major axis. This allows a kernel to be constructed in the image domain and transformed into the Fourier domain.

Our final images combine deconvolved emission (i.e., the sum of all clean components) and residuals, which are mostly noise. The deconvolved emission has been convolved with a Gaussian “clean” beam calculated from fitting the core synthesized beam. The residual emission, including any actual emission too faint to be deconvolved, still incorporates the dirty beam. Note that, by convolving the cube to have a round beam, we also change the “dirty” beam associated with this residual emission. The synthesized dirty beam can have a complex shape, but the core is similar to the elliptical Gaussian restoring beam. Therefore, the convolution to a round beam will also “round” the core of the dirty beam, thus keeping the shape of the dirty beam and restoring beam approximately similar and making the dirty beam more symmetric.

At this point, the images have a round beam, units of Jy beam^{-1} , and represent deconvolved images of the sky no longer tapered by the combined primary beam response.

Figure 10. Example of telluric ozone contamination in the PHANGS-ALMA total power observations and the effects of the correction algorithm. The left-hand panel shows a mean brightness temperature map of CO(2–1) from the southern half of NGC 1792. This target is one of six still affected by telluric ozone contamination after reobservations using a fixed-elevation OFF position. The map here was derived from a single execution block, i.e., a single observing session. The mean brightness temperature is calculated within $\pm 25 \text{ km s}^{-1}$ from the Doppler-shifted velocity of the ozone feature expected during that session ($V_{\text{LSRK}} = +1222 \text{ km s}^{-1}$). The arrow indicates the orientation of the contamination gradient derived from our fit (see text for details). The yellow polygonal line is the manual mask used to exclude the real CO emission of the target (which appears as the color-saturated red patch in the map) from the fit. We highlight the results of our algorithm in two regions at opposite ends of the gradient, labeled E and W. The middle and right-hand panels show CO(2–1) spectra averaged over the E and W regions: before correcting for any contamination (green, filled), after applying our method (black), and the difference between them (yellow, filled, offset vertically for the sake of clarity). The latter is the average spectrum of the ozone signal in either region. As expected, the ozone signal peaks at a velocity of $+1222 \text{ km s}^{-1}$, indicated with a vertical dashed line. It shifts in sign between E and W as the airmass difference between the source and the OFF shifts sign.

6.3. Stitching Together Multipart Galaxies via Linear Mosaicking

Due to the combination of ALMA’s fields-per-scheduling block limitation and its powerful mosaicking capabilities, many science projects now observe multiple large mosaics, which are then combined into a single large image during processing. In practice, for PHANGS-ALMA, many targets were observed using two or three separate mosaics. In one case, NGC 0253, five distinct maximum-sized mosaics were used. Figure 11 shows an example of a three-part observation targeting NGC 2903.

As mentioned previously, these mosaics tend to be observed at different times with different $u-v$ coverage and weather conditions, and so have different synthesized beams. CASA does not currently track positional variations in the synthesized beam, making it challenging to image all of these data simultaneously. Instead, we stitch these images together in the image domain via linear mosaicking.

To do this, we first identify a common spatial resolution for all parts of the mosaic. This common resolution is slightly larger than the coarsest resolution for any single part. As with the convolution to a round beam, this process used the CASA task `imsmooth` and required a small amount of “padding,” i.e., increasing the size of the beam in order to allow the routine to successfully carry out the convolution. This is another step where we lose a modest amount of resolution, $<10\%$. This beam matching is not an issue for the single-dish data, because the beam shape is constant.

After this convolution, we constructed a new astrometric grid that covered all of the individual mosaic parts. The individual parts already shared the same spectral axis, thanks to the pre-processing of the visibility data (Section 3). We then used the CASA task `imregrid` to align all of the parts of the mosaic to this common astrometric grid. We also aligned the combined primary beam coverage cubes onto the same grid. Last, we

aligned the single-dish total power data onto this grid for use in feathering (Section 6.4).

After this, we combined all mosaic parts into a single image. To do this, we weight each cube by the local value of the primary beam response squared times the inverse of the typical noise in that cube. That is, for each voxel, we calculate

$$\langle I \rangle = \frac{\sum_{i=0}^n (B_i^2 \sigma_i^{-2}) \times I_i}{\sum_{i=0}^n (B_i^2 \sigma_i^{-2})}. \quad (2)$$

Here, $\langle I \rangle$ is the mean intensity, which is placed in the new cube, and the sum over i refers to a sum over all mosaic parts that contribute at that voxel. The combined primary beam response squared, B_i^2 , should track regional variations of the noise within the primary-beam-corrected cube. The quantity σ_i^2 is the overall noise variance in the cube, such that the product σ_i/B_i corresponds to the local rms noise. Weighting the intensities by the inverse of their noise variance should produce the lowest possible noise in the output cube.

We did experiment with joint imaging of the mosaic parts, and found that variations in the $u-v$ coverage across the mosaic sometimes lead to divergence in the deconvolution. Stitching via linear mosaicking proved to be a much more stable option.

6.4. Combination of Total Power and Interferometric Data via Feathering

We combined the cleaned 7 m and 12 m+7 m cubes with the single-dish cubes using CASA’s `feather` task. Feather combines the interferometric and total power cubes in the Fourier domain, using the total power data at low angular frequencies (i.e., to fill in short and zero spacings) and the interferometer data for information at high angular frequencies.

Figure 11. Example of linear mosaicking for PHANGS-ALMA CO(2–1) data from the ACA 7 m antennas for NGC 2903. This galaxy was independently observed using three maximum-sized 150 field mosaics (for the 12 m array), labeled parts 1, 2, and 3. Each part is imaged separately. The individual parts are then convolved to a common beam and aligned on a shared astrometric grid. The first three panels show peak intensity images of these three aligned, beam-matched parts. The gray contours show the footprint of the individual parts. The three cubes are then combined, weighting by the local combined primary beam response and the overall noise level in the cube (Equation (2)). The final panel shows the peak intensity map from the resulting combined mosaic. We apply a similar procedure to the single-dish data, and then combine the interferometric and total power data after this mosaicking step.

We used this task with the default options, as we already reprojected the different input images to the same grid, and we ensured that they all were converted to Jy beam^{-1} units.

Figures 12 and 13 show the impact of the short-spacing correction for the PHANGS-ALMA CO(2–1) data. Figure 12 illustrates the data before and after feathering and for different arrays. We show the integrated intensity obtained from collapsing the same twenty-channel-thick slab of NGC 4303 seen in Figures 4 and 5. All six panels show the same high stretch. Also note that here we include images made using only the 12 m array and total power data. The 12 m only and 7 m

only images illustrate negative artifacts, or “bowling,” around bright emission, due to missing short-spacing data. The images that include total power data show how the single-dish data fill in the bowls and also add an extended, faint component to the image; see Pety et al. (2013) for a much more detailed demonstration in M51.

Figure 13 shows results separately for the 7 m only data and the combined 12 m+7 m data. On average, the 7 m only cubes recover $\sim 70\%$ of the emission found in the final, short-spacing-corrected cubes. The 12 m+7 m cubes do much better, recovering $\sim 90\%$ of the flux seen in the total power data on

Figure 12. Comparison of images from all array combinations for one galaxy. Integrated emission from twenty-channel-thick slabs in the CO(2–1) data cubes for the same target shown in Figures 4 and 5. Each panel shows the image for a different array combination: (top left) the jointly imaged 12 m+7 m data with no total power data; (top right) the jointly imaged 12 m+7 m data feathered with the total power data; (middle left) the 12 m data only; (middle right) the 7 m data feathered with the total power data; (bottom left) the 7 m data only; and (bottom right) the total power data only. All panels show the same high stretch, in units of K km s^{-1} , and use the same astrometric grid. The images demonstrate the artifacts, especially negative “bowls,” due to missing short-spacing data and imperfect deconvolution in 7 m only or 12 m only imaging. Contrasting the images shows how the inclusion of 7 m data improves the recovery of extended emission, compared to the 12 m only image. In Appendix C and Section 8.3, we show and discuss how the combined 12 m+7 m deconvolution recovers more flux and yields better overall results than the 7 m only case. Including the total power data almost entirely removes the negative artifacts and adds a faint, extended component in both the 12 m+7 m and 7 m cases.

average. There is a large scatter in the recovery fraction of the ACA 7 m data, with the 16–84th range spanning from 61% to 85% recovery. For high-brightness targets, the 12 m+7 m data cluster in the range 85%–99% recovery. The difference between the two arrays suggests that a large part of the

“missing flux” in the 7 m only case reflects shortcomings of the deconvolution, not only spatial filtering. We discuss the point more in Appendix C.

Other approaches: We experimented with other approaches, including `tp2vis` (Koda et al. 2019) and the use of either the

Figure 13. Fraction of PHANGS-ALMA CO(2-1) flux recovered by the interferometer before short-spacing correction. The fraction of the total flux recovered by the interferometer only for the PHANGS-ALMA CO(2-1) data, for the ACA 7 m array only (left panel) and the combined 12 m+7 m arrays (right panel). The y-axis shows the ratio between the flux in the interferometer-only integrated intensity map and the integrated intensity map constructed from the feathered data. Specifically, we calculate these fluxes from the integrated intensity maps constructed using the “broad” masks described in Section 7. The error bars combine the statistical uncertainty from noise with typical uncertainties for ALMA’s interferometric and total power calibration, added in quadrature. On average, the ACA 7 m only data recover 72% of the CO(2-1) flux seen by the total power antennas (horizontal dashed line), with the 16th–84th percentile range of 61%–85% (shaded region) across the sample. The combined 12 m+7 m imaging does much better, recovering a median 91% of the total flux with a 16th–84th percentile range of 85%–99%. The high values in low-brightness galaxies reflect low signal-to-noise cases where the imaging fails to properly deconvolve the galaxy. We discuss the better performance of the 12 m+7 m imaging compared to the 7 m only in Appendix C.

total power data or previous rounds of high S/N imaging attempts as a model (for more details, see Appendix D). To evaluate the competing methods, we created a set of images with known flux, based on collapsed versions of our CO data cubes. Next, we simulated interferometric observations of the known source using *simalma*. We simulated total power observations by convolving the true image to the resolution of the single-dish data. Then, we applied each method of reconstruction: feathering, *tp2vis*, and seeding the deconvolution using an input model.

The calculations in Appendix D yield *much* more spatial filtering than our real data. Based on comparison to the more realistic simulations carried out in Section 8.3, this appears to reflect that our actual imaging operates in individual velocity channels. As emphasized in that section, the calculations in Appendix D should be taken as experiments that consider a “worst case” scenario for spatial filtering in nearby galaxies.

These tests showed that *feather* recovered a known input image with about the same fidelity and flux accuracy as the other approaches. Typically, all of the methods implied 10%–15% inaccuracies in overall recovery of the input image, but also treat an extreme case. Still, this part of the calculation certainly represents one of the dominant uncertainties in high S/N 7 m only observations. Short-spacing correction is an area where we expect research and development to improve our data products in the coming years.

Apodization: In the current version of the pipeline, we do *not* apodize (“taper”) the single-dish image before feathering. We feather the best-estimate image of intensity from both the single-dish and interferometric data. This approach effectively treats the total power information as zero outside the field of view of the interferometer.

In theory, apodizing the single-dish and feathering before primary beam correction may seem preferable, and some CASA

and ALMA documentation recommends this approach because it carefully matches the fields of view of the two data sets and avoids any sharp edges. We conducted tests using simulated sources and emission near the edge of the field of view and found that apodizing before feathering led to distortions at the edges of the output. The leading hypothesis for this effect is that apodization interacts with the primary beam of the single-dish telescope to distort the shape of bright sources in the mosaic during feathering (C.D. Wilson et al. 2021, private communication).

Given that there is some uncertainty regarding the treatment of edges in *feather*, we carry out linear mosaicking of both the total power and interferometric data *before* feathering. In PHANGS-ALMA, most cases with bright emission near the edge of the observed field of view are part of a larger, multipart mosaic. By stitching these parts together before feathering, we minimize the impact of our treatment of the map edges.

6.5. Downsampling and Trimming of Data Cubes

After imaging and convolution to a round beam, our cubes usually have $\gtrsim 7$ pixels across the FWHM of the synthesized beam. The imaging and processing also left the cubes with a large amount of empty space surrounding the data. Both the oversampling and the padding are useful for imaging but unnecessary for scientific analysis. They also substantially inflate the data volume of the cubes. Therefore, at this stage, we trim and downsample the cubes in order to lower their volume without reducing information content.

For any cube with pixel scale fine enough that >6 pixels fit across the (now round) beam FWHM, we rebinned the cube. This rebinning increased the pixel size by a linear factor of two, which corresponded to a factor of four decrease in the number of pixels in the cube. After this rebinning, the pixels still critically sampled the beam.

Finally, we extracted only the part of each cube that contained data, dropping any extra padding in R.A. and/ or decl.

6.6. Conversion to Kelvin Intensity Scale

Finally, we convert our cubes from units of Jy beam⁻¹ to brightness temperature, T_b , measured in kelvin. This removes the beam from the units and recasts the maps onto a straightforward intensity scale, which is ideal for studying complex, resolved CO(2–1) emission.

To convert, we use the current synthesized beam size and the observed frequency in the central channel of the cube, to set a constant scaling factor. Formally, this conversion varies across our bandpass by a factor of $2\Delta\nu/\nu$, which is ~ 0.007 for the maximum 1000 km s⁻¹ bandwidth of the PHANGS–ALMA CO(2–1) cubes, but we do not include this variation. We recorded the Jansky-to-Kelvin conversion in the header of the final cube. For the PHANGS–ALMA CO(2–1) cubes, the values have 16th–84th percentile range of 0.33–0.47 K Jy⁻¹ for the ACA 7 m data and 6.2–7.0 K Jy⁻¹ for the 12 m+7 m combined data.

6.7. Exporting to FITS

At the end of this process, we export the trimmed, corrected line cubes (and images) to FITS format. During this step, we have ensured that the headers are correct and contain no extraneous information. At this stage, we have primary-beam-corrected, short-spacing-corrected, round beam data cubes in units of brightness temperature.

7. Data Product Creation

We create a series of data products from the science-ready data cubes. First, we convolve the cubes to a set of fixed angular and physical resolutions. Table 3 lists the target resolutions and other details of the product creation process. These fixed-resolution cubes are intended to allow rigorous comparison among targets at different distances (e.g., Hughes et al. 2013; Rosolowsky et al. 2021). The processing described in Section 3 already places the cubes at nearly matched velocity resolution.

For each cube, we estimate the noise at each location in the data cube. We combine this noise estimate with the data themselves to create two kinds of masks (Table 4): a “broad” mask focusing on high completeness, and a “strict” mask focusing on including only emission detected at high confidence.

We apply these masks and collapse the cubes along the spectral axis to produce a variety of “moment” maps (Table 5): integrated intensity, peak intensity, intensity-weighted mean velocity, line width, and so on. Whenever feasible, we also calculate corresponding uncertainty maps.

In the PHANGS–ALMA pipeline, these operations occur outside of CASA, in a Python environment. We use routines built around the `numpy`, `scipy`, `astropy`, `radio-beam`, and `spectral-cube` packages.

For continuum products, we carry out the convolution and estimate a single noise value from the signal-free region of the image. Most of this section describes processing of data cubes.

Table 3
PHANGS–ALMA Resolution and Noise

Item	Description
Resolutions^a	
7 m+TP galaxies	84 galaxies ^b
7 m+TP native angular ["]	7.6 ^{+0.8} _{-0.5}
7 m+TP native physical ^c [pc]	550 ⁺¹⁴⁰ ₋₁₅₀
12 m+7 m+TP galaxies	77 galaxies ^b
12 m+7 m+TP native angular ["]	1.3 ^{+0.4} _{-0.2}
12 m+7 m+TP native physical ^c [pc]	100 ⁺³¹ ₋₃₅
Common resolutions (when allowed by data)	
Angular	native, 2", 7"5, 11", 15"
Physical ^c [pc]	60, 90, 120, 150, 500, 750, 1000
Noise in individual 2.54 km s⁻¹ channels	
7 m+TP galaxies	84 galaxies ^b
7 m+TP median noise, native res. [mK]	12 ^{+4.5} _{-3.5}
7 m+TP median noise, 750 pc [mK]	7.1 ^{+3.8} _{-3.6}
7 m+TP full fractional spectral variation	0.25 ^{+0.06} _{-0.06}
7 m+TP $\pm 1\sigma$ fractional spatial variation	0.80 ^{+0.10} _{-0.17}
12 m+7 m+TP galaxies	77 galaxies ^b
12 m+7 m+TP median value, native res. [mK]	85 ⁺⁴⁰ ₋₄₀
12 m+7 m+TP median value, 150 pc [mK]	53 ⁺²⁵ ₋₂₀
12 m+7 m+TP full fractional spectral variation	0.23 ^{+0.05} _{-0.09}
12 m+7 m+TP $\pm 1\sigma$ fractional spatial variation	1.0 ^{+0.2} _{-0.28}

Notes. These numbers refer to the first public data release, internal “version 4,” constructed with “version 2.0” of the PHANGS–ALMA pipeline. They refer to the products created for the CO(2–1) survey. The number of galaxies indicates the number of targets with these array combinations processed by this release. The $\pm$ values refer to the 16th and 84th percentiles of the sample distribution.

^a Data are convolved to each of these resolutions whenever the native resolution is fine enough to allow this. Target resolutions are the same for configurations with and without total power. The quoted value refers to the FWHM of a Gaussian beam.

^b When this table was compiled, six galaxies were still missing total power data, due to the telluric contamination described in Section 5. Since then, these data have been corrected and all 90 galaxies have total power data. The median properties of the data are essentially unchanged. In total, in the public data release, 81 galaxies have 12 m+7 m+TP data and 90 have 7 m+TP.

^c When convolving to a fixed physical resolution, we adopt the current best estimate of the galaxy’s distance (from Anand et al. (2021), for PHANGS–ALMA).

7.1. Convolution to Fixed Resolutions

We convolve each data cube to a series of fixed angular and physical resolutions. This has two purposes. First, convolving to coarser angular resolution improves the surface brightness sensitivity and increases the fraction of the flux detected at good S/N (see Figure 14). This allows us to use coarser resolution versions of the cube to create high-completeness masks to be applied to the sharper resolution data. Second, convolving to multiple fixed physical resolutions plays a crucial role in testing scientific hypotheses. At the most basic level, this allows for rigorous comparison among galaxies observed with different beams and lying at different distances (e.g., Hughes et al. 2013; Rosolowsky et al. 2021). Increasingly, spatial scale is also by itself viewed as an important

Table 4
PHANGS–ALMA Masking

Item	Description
Mask summary	
Strict mask	Low false-positive rate
Strict mask	Based on S/N threshold
Strict mask	Constructed for every resolution
Broad mask	High completeness
Broad mask	Union of strict masks for all resolutions
Broad mask	One mask per configuration and target
Completeness ^a	
7 m+TP galaxies	84 galaxies
7 m+TP strict mask-to-direct sum	$0.77^{+0.11}_{-0.12}$
7 m+TP broad mask-to-direct sum	$0.92^{+0.06}_{-0.12}$
7 m+TP strict mask-to-broad mask	$0.84^{+0.08}_{-0.14}$
12 m+7 m+TP galaxies	77 galaxies
12 m+7 m+TP strict mask-to-direct sum	$0.62^{+0.20}_{-0.25}$
12 m+7 m+TP broad mask-to-direct sum	$0.98^{+0.08}_{-0.07}$
12 m+7 m+TP strict mask-to-broad mask	$0.65^{+0.17}_{-0.25}$

Notes. These numbers refer to the first public data release, internal “version 4,” constructed with “version 2.0” of the PHANGS–ALMA pipeline. They refer to the products created for the CO(2–1) survey. The number of galaxies indicates the number of targets with these array combinations processed by this release.

^a Completeness here refers to the fraction of flux included in each mask at the native resolution of the cube. For both masks, we reference this to a direct sum of the cube, also at the native resolution. We also calculate the ratio of flux between the two masks. Quoted values are medians, and the error bars refer to the 16th and 84th percentile. These measurements are visualized in Figure 20.

variable when studying stochastic processes and the hierarchical structure of the interstellar medium (e.g., Schrubba et al. 2010; Schinnerer et al. 2019; Chevance et al. 2020).

For PHANGS–ALMA, we convolved the data to a series of fixed angular and physical scales. The target angular scales have FWHM beam sizes of 2″, 7″.5, 11″, and 15″, and the fixed physical resolutions are 60, 90, 120, 150, 500, 750, and 1000 pc. When convolving to a fixed physical resolution, we adopt a distance to the galaxy, which the user supplies as an input to the pipeline. We then calculate the angular scale corresponding to the target physical resolution at the adopted distance of the galaxy. We use Euclidean geometry for this calculation. For PHANGS–ALMA, we adopt the distances derived and compiled by Anand et al. (2021). For these target angular resolutions, the 12 m+7 m data can typically be convolved to all scales. The 7 m data can all be convolved to 11″ and 15″, but less than half can be convolved to 7″.5. Roughly speaking, the target resolutions of 60, 90, 120, and 150 pc correspond to the quartiles of the distribution of 12 m+7 m physical resolutions (see Figure 15 and Table 3). For a typical target, the 7 m data can only reach $\gtrsim 500$ pc resolution, but a key extension of PHANGS–ALMA targets galaxies with $d < 5$ Mpc. In these targets, the 7 m data can also reach physical resolutions $\lesssim 165$ pc.

When we convolve the cubes, we treat any area outside the map as missing. This means that, near the edges of the map, comparatively fewer data contribute to the final map. As a result, the noise will be higher near the map edges. We create a “coverage cube” to track the amount of data contributing to each

sight line. To create this cube, we replace all locations with data in the original cube with 1.0 and all locations without data with 0.0. Then, we also convolve this cube. In the resulting coverage cube, a value of 0.95 indicates that, during the convolution, 95% of the effective area of the convolving beam contained data, i.e., had values of 1.0, while 5% of the convolved area did not, i.e., had values of 0.0. We use this coverage cube to clip some final data products to avoid strong edge effects.

Application to PHANGS–ALMA: Table 3 and Figures 14 and 15 illustrate some details of the resolution and convolution for PHANGS–ALMA. Table 3 and Figure 15 report the native angular and physical resolutions of the 12 m+7 m+TP and 7 m+TP data after postprocessing. The 7 m+TP data show a narrow range of angular resolutions, consistent with the almost fixed configuration used to observe them. The distances to nearby galaxies, including the PHANGS–ALMA targets, are almost always uncertain by 5%–30%; for more details, see, e.g., Tully et al. (2016), McQuinn et al. (2017), and Anand et al. (2021), among many others. After accounting for the current best-estimate distances to the targets (Anand et al. 2021), the 7 m+TP data show a wide range of physical resolutions, typically ~ 550 pc but with outliers down to < 200 pc. These high resolutions arise from 7 m observations of very nearby systems with $d \lesssim 5$ Mpc, many of which we have so far targeted only with the ACA.

In Figure 15 and Table 3, the 12 m+7 m+TP data show a wider range of angular resolutions. This mostly reflects that ALMA delivers data within some tolerance of the nominal angular resolution and that the 12 m array cycles between array configurations. As a result, the exact $u-v$ coverage differs from galaxy to galaxy. After accounting for distance, the typical physical resolution of the 12 m+7 m+TP data is 100 pc, with the highest resolution ~ 25 pc, all galaxies better than 200 pc, and 90% of galaxies having physical resolution better than 150 pc.

Note that the resolutions in these final cubes have been inflated by several postprocessing steps (Section 6). We convolved to a round synthesized beam and also degraded to the coarsest common resolution when linearly mosaicking individual “parts” of multipart mosaic galaxies. Each of these steps involves a convolution to a moderately coarser resolution, and in the case of multipart galaxies, one part may have much higher resolution than the other (a prominent example of this in PHANGS–ALMA is NGC 4321, M100, where one half of the galaxy has much higher resolution (1″.0) than the other (1″.6); see, e.g., Henshaw et al. (2020)).

Figure 14 shows an example of the convolution to fixed physical resolution applied to one PHANGS–ALMA galaxy, NGC 3621. Each panel shows the galaxy at a fixed physical resolution, from 60 pc to 1 kpc. The contour shows the area of high ($>95\%$) coverage as defined above. The figure shows increased surface brightness sensitivity, increased filling fraction of emission, and decreased detail as the resolution degrades. The bottom right panel also illustrates how, at coarser resolutions, edge effects become important. As the beam becomes larger, a larger fraction of the original flux sits near the edge of the field of view, where the sensitivity is reduced due to incomplete sampling.

7.2. Noise Estimation

For each cube at each spatial scale, we produce a three-dimensional estimate of the rms noise. We treat this as a separable problem. First, we construct a noise map that

Table 5
PHANGS–ALMA Derived Product Summary

Map	Expression	Unit	Uncertainty Method
Integrated intensity (mom0)	$W(x, y) = \sum_i I(x, y, v_i) M(x, y, v_i) \delta v$	K km s^{-1}	Gaussian
Peak intensity (tpeak)	$I_{\text{peak}}(x, y) = \max_{v_i} [I(x, y, v_i)]$	K	None
Peak intensity (smoothed) ^a (tpeaklp5)	$I_{\text{peak}, \Delta V}(x, y) = \max_{v_i} [I(x, y, v_i) * K(\Delta V)]$	K	None
Mean velocity (mom1)	$\bar{v}(x, y) = \frac{1}{W(x, y)} \sum_i v_i I(x, y, v_i) M(x, y, v_i) \delta v$	km s^{-1}	Gaussian
Velocity at peak intensity (vpeak)	$v_{\text{peak}}(x, y) = \text{argmax}_{v_i} [I(x, y, v_i)]$	km s^{-1}	None
Interpolated peak velocity (vquad)	$v_{\text{quad}}(x, y) = v_{\text{peak}}(x, y) - \frac{A}{B}$ for $A = I(x, y, v_{\text{peak}} + 1) - I(x, y, v_{\text{peak}} - 1)$ $B = I(x, y, v_{\text{peak}} - 1) + I(x, y, v_{\text{peak}} + 1) - 2I(x, y, v_{\text{peak}})$	km s^{-1}	Gaussian
rms line width ...(mom2)	$\sigma_v(x, y) = \left[\frac{1}{W(x, y)} \sum_i (v_i - \bar{v})^2 I(x, y, v_i) M(x, y, v_i) \delta v \right]^{1/2}$	km s^{-1}	Gaussian
Equivalent/effective width ...(ew)	$\text{EW}(x, y) = [\sum_i I(x, y, v_i) \delta v] / [\sqrt{2\pi} I_{\text{peak}}(x, y)]$	km s^{-1}	Gaussian

Notes. For all entries, $I(x, y, v)$ is the position–position–velocity data cube produced by the pipeline, $M(x, y, v)$ is a Boolean mask indicating where CO emission is found, and δv is the channel width, taken as a constant.

^a Here, $K(\Delta V)$ is a boxcar smoothing kernel of full width $\Delta V = 12.5 \text{ km s}^{-1}$ and $*$ is the convolution operator.

captures spatial variations of the noise, $R(x, y)$. Then, we measure a normalized noise spectrum that captures the relative spectral variations, $s(v)$. The noise in the data cube is then:

$$\sigma(x, y, v) = R(x, y) s(v). \quad (3)$$

We determine $R(x, y)$ and $s(v)$ empirically, determining the values from the data themselves using an iterative procedure. First, we use a robust noise estimator, the median absolute deviation of the data around zero, to characterize the noise in a rolling spatial box. In estimating the noise, we exclude positive data that have high significance with respect to the noise level. These are likely to be associated with real emission and not noise. We do not exclude high-significance negative data, because they have not been a concern for PHANGS–ALMA, but we might modify the calculation to do this in the future.

To save computation time and increase the sample size of the data used for noise estimation, we calculate the noise in boxes centered on a sparsely sampled square grid rather than at every pixel in the cube. The size of the box and the grid spacing are tunable parameters. Larger boxes yield more robust noise estimates, thanks to the large sample size, at the expense of washing out small-scale variations in the noise. For the PHANGS–ALMA public release, we used a box size of $\sim 3 \times$ the FWHM beam size in width. We calculate the noise on a rectilinear grid of positions with spacing of 1.2 beam FWHM. We then smooth the empirical noise estimates with a Gaussian kernel with size equal to the box size, yielding an estimate of $R(x, y)$.

We then estimate $s(v)$ by normalizing the cube by the spatial response, R , and estimating the median factor that each channel is different from the spatial noise estimate derived for the cube as a whole. We then smooth these estimates with a third-order Savitsky–Golay filter, to estimate $s(v)$. In PHANGS–ALMA, both the performance of the receiver and the regridding effects described in Section 3 lead to spectral variations of the noise.

This process is iterative. We generate an estimate of $\sigma(x, y, v)$, divide the cube by this estimate, and then repeat the noise

estimation process. Variations in the noise estimate are accumulated to form a final estimate of $\sigma(x, y, v)$. The iterative process drives $I(x, y, v)/\sigma(x, y, v)$ in the signal-free regions to a zero-centered normal distribution with standard deviation of 1. In practice, we find that three iterations are sufficient to arrive at a stable estimate of the noise. Figure 16 shows the variations seen in a typical noise map and that the resulting noise cube characterizes the spatial and spectral variations of noise in the cube.

Application to PHANGS–ALMA: Table 3 and Figures 16–18 report some results of applying this algorithm to the PHANGS–ALMA CO(2–1) data. Table 3 and Figures 17 and 18 report typical noise values and typical spatial and spectral variations in the cubes. We show the normalized noise spectra of each galaxy in Figure 17.

Table 3 and Figure 18 show median noise levels of 12 mK for the native resolution 7 m+TP data, 7 mK for the 750 pc resolution 7 m+TP data, 85 mK for the native-resolution 12 m+7 m+TP data, and 53 mK for the 150 pc resolution 12 m+7 m+TP data. In each case, individual galaxies scatter by $\sim \pm 50\%$ about these median values. As expected, the convolution lowers the overall noise level, but the fractional scatter in the data set remains about the same at each resolution.

Table 5 notes the magnitude of spatial and spectral noise variation across the final PHANGS–ALMA cubes. We typically find $\sim 25\%$ variation in the spectral dimension. We observe much larger spatial variations, with $\pm 1\sigma$ variations of 80%–100% on average (i.e., the 84th–16th percentile value divided by the median is ~ 0.8 –1.0). As shown in Figure 16, this mostly reflects the large variation of noise near the map edge, due to the changing primary beam response. For galaxies observed in multiple parts (e.g., Figure 11), the different parts often have different surface brightness sensitivities. This also contributes to the spatial variation.

For most galaxies, the noise spectra shown in Figure 17 exhibit a common behavior. The two arrays also show similar behavior to one another. The noise tends to increase from low

Figure 14. Example of data products and coverage derived from the convolved cubes. This figure shows peak temperature maps using a 12.5 km s^{-1} spectral window at four resolutions for NGC 3621. Each panel shows the product derived after convolution to a different physical resolution (from top left to bottom right): 60, 150, 500, and 1000 pc. The image is set to cover the 1%–99% range of the data on an arcsinh stretch. The contours show 5, 10, 20, 40, 80, and 160 times the noise level. Circles in the lower left of each panel show the smoothed beam FWHM sizes. The images show that, as the resolution degrades, the sensitivity and extent of detections increases but fine details are washed out. Blue lines show the area of 95% coverage after the convolution. Regions outside this contour include fewer data than those inside, and so suffer from increasing edge effects as one approaches the map edge.

to high recessional velocity, i.e., with decreasing frequency. The decrease has a coherent shape across most galaxies, with median variation magnitude of $\sim 25\%$ in both arrays. We understand this overall gradient as a combined result of the spectral regridding effects discussed in Section 3.4 and the behavior of the Band 6 receiver used to make the measurement. The spectral gridding introduces a gradual gradient across the bandpass, as slightly different amounts of independent data contribute to different channels. The receiver effect refers to the fact that we place the CO line relatively close to the lower edge of the upper sideband of the ALMA Band 6 receiver. The receiver temperature rises with decreasing frequency in this regime (C. Brogan et al. 2021, private communication). One notable outlier in the $12 \text{ m} + 7 \text{ m} + \text{TP}$ plot is NGC 0628, where we placed the line in the middle of the lower sideband.

We see the same trend in noise as a function of velocity in the $7 \text{ m} + \text{TP}$ data, but we also find enhanced noise near the systemic velocity of the galaxy. This reflects the fact that our

iterative noise rejection does not do a perfect job of filtering out the emission from the galaxy in this case. The emission in the $7 \text{ m} + \text{TP}$ maps tends to be more extended, with a larger filling factor and higher median S/N as compared to the $12 \text{ m} + 7 \text{ m} + \text{TP}$ maps. As a result, it appears to bias our noise estimates high by about 10% over the velocity range of the galaxy. Other than this effect, the average spectral variation of the noise matches well between the $12 \text{ m} + 7 \text{ m} + \text{TP}$ and $7 \text{ m} + \text{TP}$ data. That is, the blue and red lines overlap away from the systemic velocity in the two panels of Figure 17.

Finally, recall that, at several steps during the imaging (Section 4), we use a single robustly determined noise value to describe the data, rather than the three-dimensional estimate here. Note that this processing happens before any primary beam correction, and it treats individual mosaics separately. Therefore, most of the spatial noise variations will be suppressed. We expect these estimates to be accurate to $\sim 30\%$.

Figure 15. Native physical and angular resolutions of the PHANGS-ALMA cubes. The panels show histograms of the native angular (top row) and physical (bottom row) resolutions of the PHANGS-ALMA 7 m+TP (left) and 12 m+7 m+TP (right) cubes after processing (Table 5). To calculate the physical resolutions, we adopt the distance compilation from Anand et al. (2021). The wider range of angular resolutions for the 12 m+7 m data reflects the fact that the 12 m array configuration used to take these data varied somewhat, consistent with standard ALMA observing strategies. Extreme outliers at fine physical resolution for the 7 m+TP data reflect cases where we have used the ACA 7 m+TP only to target very nearby, extended systems.

7.3. Masking

Since the data cubes include large, signal-free volumes, we create masks to identify the regions of the cube containing signal. We then apply these when creating higher-level data products.

We create two types of masks, which we illustrate in Figure 19. First, we create a high-confidence “strict mask” that includes only voxels highly likely to contain real signal. Second, we create a high-completeness “broad mask,” which contains most known signal in the cube. Though there are many approaches to masking, these two cases cover most common applications. The strict mask, which is illustrated in the right column of Figure 19, includes only bright emission and few or no noise-dominated sight lines. It should be used when running calculations sensitive to noise, e.g., many types of kinematic analysis. The broad mask, illustrated in the left column of Figure 19, should include almost all regions with real emission. This comes at the expense of including more noise-dominated sight lines. The broad mask should be used for any analysis aimed at a complete characterization of the emission.

We create strict masks for each cube at each resolution. These mostly follow the standard recipes defined for CPROPS (Rosolowsky & Leroy 2006). They begin with a core mask that includes all voxels with S/N above 4 over two successive velocity channels.⁴⁹ We also create a lower S/N outer mask that includes all voxels with S/N above 2 in two successive velocity channels. We then construct a final mask that consists of all contiguous regions in the outer mask that contain any pixels from the higher-significance core mask. As long as the channel width is a few times narrower than the typical line width, this algorithm does an excellent job of identifying all

significant features in the cube. The example in the right column of Figure 19 shows that, for NGC 4303, the strict mask indeed does a good job of highlighting all of the real emission one would pick out from a peak temperature map.

We offer the user the option to trim small-volume or small-area regions from the core mask, with both specified in units of the beam area. In order to maintain a relatively clean, easily modeled criteria for inclusion in the mask that applies for individual lines of sight (Sun et al. 2018, 2020), we do not use these volume options for the main PHANGS-ALMA data products. We do apply one additional condition on the strict mask, however. When masking data cubes that have been created by convolution, we restrict the core mask to only include regions that had high coverage in the original map. Specifically, we only allow regions that have a value greater than 0.95 in the “coverage cube” (see above) to contribute to the core mask. This prevents spurious contributions from map edges, where the noise estimate can become slightly inaccurate.

We create broad masks by taking the union of all strict masks from all resolutions. Both high-resolution and low-resolution masks contribute to the final result. The masks at the coarse resolution tend to do an excellent job of capturing extended, faint emission. These tend to be most important for overall recovery of flux in PHANGS-ALMA targets. The masks at high resolution tend to capture bright compact features, e.g., these high-resolution masks do a better job of recovering the broad line wings associated with galactic nuclei than do the low-resolution ones. By combining these masks, we construct a best estimate of where we have detected any signal in the cube at any resolution. As illustrated in the left column of Figure 19, the broad masks do a good job of encompassing all emission from the galaxy, at the expense of including a moderate amount of “empty” noise-dominated volume. In many PHANGS-ALMA cases, including the one illustrated, the broad mask captures the overall rotation of the galaxy and extends across most of the area of the map.

Note that the broad masks resemble the clean masks described in Section 4. For PHANGS-ALMA, these are not

⁴⁹ In principle, the appropriate number of channels required for joint detection depends on the expected line width and the line spread function. In PHANGS-ALMA, two channels correspond to a line width of $\sim 5 \text{ km s}^{-1}$, about the narrowest full line width that we might expect for a GMC (e.g., Bolatto et al. 2008; Heyer et al. 2009). The channels are also mostly independent (when tested following the method in Leroy et al. (2016)), such that the line spread function is about one channel.

Figure 16. Examples of empirically generated noise cubes. The top and middle panels show the noise map in the spatial (top) and spectral (middle) coordinates for the 12 m+7 m imaging of NGC 4303. Contours in the top panel show the 16th (blue), 50th (green), and 84th (blue) percentiles of the noise values. The noise profile along the spectral axis is extracted from the center of the map. The bottom panel shows the probability density function of the signal-to-noise implied by this noise cube. The blue parabola shows the PDF of a normal distribution with mean of zero and variance of one. A normal distribution is an excellent description of the signal-to-noise values, except for the strong positive tail of values arising from signal in the cube.

identical. We create external clean masks and supply them. If one wanted to use the PHANGS–ALMA pipeline to create clean masks in an automated way, one could process the data, create broad masks, then feed them back in as clean masks. In practice, the main differences between our broad and clean masks are that the clean masks had an additional dilation in all three dimensions (i.e., they have been slightly “inflated”) and that the clean masks for galaxies with bright centers include a wide velocity region near the center of the galaxy (compare Figures 6 and 19).

Application to PHANGS–ALMA: Table 4 and Figures 19 and 20 show some outcomes of applying this masking to PHANGS–ALMA. Figure 19 illustrates the differences between the strict and broad masks for a typical bright galaxy, NGC 4303. Table 4 and Figure 20 report the fraction of flux captured by each mask for our data.

For the 12 m+7 m+TP data, the strict masks have $\sim 60\%$ completeness on average, meaning that they include about 60% of the emission found via a direct sum of the cube. As in Sun et al. (2018, 2020), we find a wide range of completeness among the PHANGS–ALMA data, with the 16%–84% range spanning about 40%–80%. There is not a perfect mapping between integrated CO flux and completeness in the strict maps, but our lowest-completeness galaxies do tend to have lower overall flux. These are often, but not always, lower-mass, more HI-dominated systems. For the 12 m+7 m+TP data, with only a few exceptions, the broad masks do a good job of achieving nearly 100% completeness. The outliers tend to be the lowest-flux galaxies. If we fail to detect diffuse signal in any mask, even at low resolution, the broad mask will underestimate the true flux. Masked flux fractions larger than unity can occur because the masked regions do not include the negative noise fluctuations that are included in the sum over the cube.

For 7 m+TP data, the strict masks have higher overall completeness, almost 80% on average, with a range of about 70%–90%. The completeness of the broad mask compared to direct integration of the cube is actually moderately lower for the 7 m+TP data than the 12 m+7 m+TP data: only 92%, on average, for the 7 m+TP data. This likely reflects the fact that our set of spatial scales only reaches to $15''$, which is still somewhat compact compared to the 7 m native resolution of $\sim 7''.5$. Still, the completeness of the broad masks for the 7 m+TP data is quite high for all high-flux targets. As with the 12 m+7 m+TP data, the completeness drops in faint, lower surface brightness targets. Because there is less difference between the broad and the strict masks for the 7 m+TP data, the completeness of the two track one another closely as a function of total flux, but with the strict masks mildly offset to lower completeness.

Overall, the masks perform as intended in PHANGS–ALMA. The broad masks achieve nearly 100% completeness in many cases, while the strict masks have lower completeness but higher confidence. The difference is much less marked in the 7 m+TP data compared to the 12 m+7 m+TP data, because the strict masks already have high completeness due to the high surface brightness sensitivity of the 7 m+TP data.

7.4. Map Creation

We combine the cubes, noise estimates, and masks to produce a suite of high-level data products and associated uncertainties. In general, we deliver each product at each possible resolution. Figures 21–23 illustrate these products for one PHANGS–ALMA galaxy.

We produce associated uncertainty maps via Gaussian error propagation. For a map that estimates a two-dimensional product over a spectrum using a function f , the variance in our estimate of f will be

$$\sigma_f^2 = \sum_{i,j} \left(\frac{\partial f}{\partial v_i} \right) \left(\frac{\partial f}{\partial v_j} \right) \sigma_{ij}^2, \quad (4)$$

where σ_{ij}^2 is the variance-covariance matrix with $\sigma_{ii}^2 = \sigma_i^2$, and the sum runs over channels i and j , with v_i and v_j being the

Figure 17. Normalized noise spectra for PHANGS-ALMA data cubes. Both panels show normalized noise, calculated by our three-dimensional noise estimator (Section 7.2) as a function of Doppler shift velocity for PHANGS-ALMA galaxies. We plot noise divided by the median noise in the cube and velocity offset from the mean velocity in the cube. We calculated the median noise in the cube from a 100 km s^{-1} wide window at each edge of the spectrum. Most galaxies and both array combinations show the same overall trend in noise as a function of velocity. We attribute this to a mixture of the gridding effects discussed in Section 3.4 and the noise response of the Band 6 receiver used for the survey (e.g., Figure 22 in Kerr et al. 2014 and C. Brogan 2021, private communication). The smooth trend in the 12 m + 7 m + TP data (left panel) suggests that the iterative signal rejection works well for these data. The 7 m + TP data (right) show the imprint of the galaxy emission superimposed on the background trend near the mean (systemic) velocity. This modest ($\sim 10\%$) effect reflects that our rejection of signal from the noise estimate works well but not perfectly in these lower-resolution cases.

Figure 18. Characteristic noise in the PHANGS-ALMA cubes. Each plot shows characteristic noise values, in units of millikelvin, for each PHANGS-ALMA 7 m + TP and 12 m + 7 m + TP cube at their native resolution (top row). These “characteristic” values are drawn from the center of the three-dimensional noise cube. Spatial (Figure 16) and spectral (Figure 17) noise estimates represent deviations about this. The lower panels show the noise after convolving the 7 m + TP data to 750 pc resolution and the 12 m + 7 m + TP data to 150 pc resolution, omitting galaxies that cannot be convolved to this resolution because of their distance and angular resolution. The convolution lowers the characteristic noise (Table 5).

intensity in each channel. We use the three-dimensional error estimates to determine the uncertainties from Section 7.2. We can also include the effects of channel-to-channel correlation in our uncertainties. We model the covariance between channels in terms of a correlation coefficient, r , measured as a function of channel separation. The covariance is then

$$\sigma_{ij}^2 = r(|i - j|) \sigma_i \sigma_j. \quad (5)$$

We measure the channel correlation empirically from our imaging products (e.g., Leroy et al. 2016; Koch et al. 2018a) and find $r(0) = 1$, $r(1) \approx 0.05$, and $r \approx 0$ otherwise. This

implies that covariance between channels increases σ_f^2 by $\sim 10\%$ relative to the uncorrelated case.

The pipeline produces the following data products as summarized in Table 5.

1. *Integrated intensity* (mom0): We integrate the cube along the spectral dimension to produce the integrated intensity in units of K km s^{-1} , also referred to as the “moment 0” map. We create versions using both the strict and broad maps. These products and the associated uncertainties (emom0) represent our basic assessment of the distribution of line emission on the sky. The “broad” versions of these

Figure 19. Illustration of our two masking schemes. The left column shows the “broad mask” in a black-and-white contour over a peak intensity map collapsed along the spectral dimension (top) and the R.A. direction (bottom) for NGC 4303. The contour shows all lines of sight where the mask has at least one “True” (1) value along the line of sight. The right panels show the same images, but contours now indicate the “strict mask.” The emission identified by the strict mask can be distinguished from noise with high confidence. This figure shows that the strict mask roughly includes all of the emission that can be seen in the peak intensity maps on a high stretch. We construct the broad mask from the union of strict masks made at all resolutions. The broad mask has high completeness, meaning that it contains most of the emission in the cube. As illustrated here, the broad mask often extends across the whole map and tracks the circular rotation of the galaxy.

maps should show the location of essentially all emission in the cube. Figure 21 shows both the broad and strict versions of these maps for NGC 4303 at 150 pc resolution.

2. *Peak intensity with and without a matched-line width filter* (t_{peak} and $t_{\text{peak}12\text{p}5\text{kms}}$): We calculate the peak intensity along each line of sight, in units of kelvin. Such “peak temperature” maps offer a useful way to see faint signal and highlight structure in the cube with minimal masking.

We also found it very useful to create “matched-line width” versions of the peak temperature map. To produce these, we smooth the data cube along the spectral dimension using a tophat kernel with width equal to the expected line width. We used 12.5 km s^{-1} for PHANGS–ALMA. These matched-filter peak intensity maps produce some of the cleanest views of faint structure in the cubes. Figure 21 shows peak temperature maps of NGC 4303 at

150 pc resolution after applying the broad mask. We show both the single-channel and 12.5 km s^{-1} wide versions.

3. *Intensity-weighted mean velocity, with and without priors, and velocity at peak intensity* (mom1 , mom1w-prior , v_{peak} , and v_{quad}): We calculate the intensity-weighted mean velocity of each spectrum, in units of km s^{-1} , also known as the “moment 1” map. We also calculate the uncertainty associated with this map (emom1). The left panel of Figure 22 shows the intensity-weighted mean velocity field calculated after applying the strict mask to the 150 pc resolution version of the NGC 4303 PHANGS–ALMA CO(2–1) cube.

We also record the velocity associated with the peak intensity and the centroid velocity near the peak calculated following Teague & Foreman-Mackey (2018). This estimator uses a quadratic function to interpolate the spectral coordinate of the local maximum

Figure 20. Flux recovery in our two masking schemes. Fraction of flux recovered using the moment 0, i.e., integrated intensity, maps created using the “strict” (red) and “broad” (blue) masks. In both panels, we show the ratio of flux in the masked moment maps to the total flux calculated by summing the entire cube. Shaded regions and lines show the 16%–84% range and median for each type of mask (see also Table 4). High-confidence “strict” masks produce moment maps that include less of the overall flux: $\sim 60\%$ on average for the 12 m+7 m+TP data, and $\sim 80\%$ on average for the 7 m+TP data. However, each sight line in a strictly masked moment map is highly likely to contain real emission (see Figure 19 and Sun et al. 2018, 2020). Maps constructed using the high-completeness broad masks include almost all flux for the 12 m+7 m+TP data and are $\gtrsim 90\%$ complete for the 7 m+TP data. A few cases show ratios above 1, which could result from mild calibration differences between the total power and interferometer data—or, more likely, failure of the direct integral of the cube to yield an accurate flux, e.g., due to mild baseline issues or field-of-view clipping effects in small maps.

at subchannel resolution. We also calculate the uncertainty from this estimator.

Furthermore, we create a version of the velocity field designed to include measurements with lower S/N than those captured by the strict mask but reject data likely to represent outliers. This map includes all moment 1 values derived from the strict mask. It also includes all moment 1 values calculated from the broad mask that meet three conditions: (i) there is no strict mask measurement for that line of sight, (ii) the integrated intensity (i.e., moment 0) value along that line of sight calculated after applying the broad mask has S/N above some threshold, and (iii) the measured velocity is within some tolerance of a prior guess at the velocity field.

For PHANGS–ALMA, this “moment 1 with prior” includes lines of sight with moment 0 S/N above 2, uses the $15''$ resolution moment 1 map as a prior, and allows values within $\pm 30 \text{ km s}^{-1}$ of that prior. In some cases, this approach can dramatically expand the coverage of the velocity field at high resolution (e.g., see Lang et al. 2020). This processing follows the approach applied by Colombo et al. (2014) to M51. They used a model rotating disk with the measured M51 rotation curve as the prior. Using the CO rotation curves of Lang et al. (2020) might be our approach in a future release.

The right panel of Figure 22 shows the intensity-weighted mean velocity field using this hybridization and prior technique. The figure shows a dramatic expansion in area covered compared to the moment 1 calculated from the strictly masked cube, but still reveals a coherent velocity structure with only modest impact from noise.

4. *Root-mean-square line width and “effective width” or “equivalent width”* (mom2or σ_v and ew): We record the intensity-weighted rms scatter of emission about the intensity-weighted mean velocity, i.e., the second

moment, and the associated uncertainty. For a Gaussian line profile, this corresponds to the 1σ width of the line. Because this estimator becomes unstable in the presence of noise, we only calculate it for the strict maps. This use of the strict maps can, in turn, bias this line width to low values, because faint line wings can be missed by the strict mask. In these cases, best practice is to correct σ_v using an analytic or data-driven extrapolation (e.g., see Rosolowsky & Leroy 2006), or to cleanly define a selection function when studying line widths (e.g., Sun et al. 2020). The pipeline does not currently implement any such clipping or sensitivity correction for σ_v .

We also record the “effective width” or “equivalent width” of each line following the definition of Heyer et al. (2001), as well as the uncertainty. This definition of line width is more robust to noise and outliers compared to the second moment, but more sensitive to the velocity resolution of the data and shape of the line profile. In this definition (see Table 5), the effective width is the integrated intensity divided by the peak intensity. Note that this differs from the optical definition of equivalent width. The name “effective width” has been suggested to avoid confusion. Figure 23 shows both the moment 2 and effective width maps for the PHANGS–ALMA map of NGC 4303 after applying the strict mask.

For quantities currently without associated uncertainties, our recommendation is to use a Monte Carlo calculation along with the noise cube to simulate uncertainties.

8. Quality Assurance and Regression Tests

To ensure that the PHANGS–ALMA data products were science-ready, we implemented a set of quality assurance (QA) and regression procedures. During the initial internal data releases, we built a detailed report for each data cube, which

Figure 21. Products showing integrated and peak intensity. The figure shows four views of intensity for NGC 4303. The top row shows the line-integrated intensity, or “moment 0,” calculated after applying the strict (left) or broad (right) mask to the data. The two distributions appear very similar, but the broad mask includes more area, including many sightlines with little or no emission. The bottom row shows peak intensity calculated using either a 2.54 km s^{-1} spectral window (left) or a 12.5 km s^{-1} spectral window (right). Both do an excellent job of highlighting faint structure, e.g., in the interarm regions. The 12.5 km s^{-1} filter used in the right-hand map approximately matches the typical line width of emission. As a result, it has moderately lower noise and higher signal-to-noise.

was passed to two experts for a careful by-eye inspection. Later in the project, we constructed automated regression tests, which we benchmarked against previous versions of the imaging. We also carried out an end-to-end check on the staging, imaging, and postprocessing using a simulated data set. This check verified the ability of the pipeline to recover a known input image. All of these tests helped to highlight several subtle issues related to the accuracy of the deconvolution, which we discuss in more detail in Appendix C.

8.1. Manual Data Inspection

For the initial internal PHANGS–ALMA data releases, we generated a collection of plots and tables that we refer to as a

“QA report” for each data cube. These reports were distributed among ~ 15 team members with experience analyzing millimeter data. Each report was assigned to at least two reviewers.

The reviewers were asked to identify potential pathologies and assess the overall quality of the image. Their feedback was used primarily to identify major failure cases, but also to improve our overall deconvolution and data processing strategies. As an example of this feedback, the spectral noise patterns discussed in Section 3.4 were first identified during an early round of QA report inspection.

Contents of the Inspection Reports: The QA report aimed to present the data in a digestible form that captured the properties of the emission, characterized the noise, and highlighted any

Figure 22. Data products showing the velocity field. The left panel shows the intensity-weighted mean velocity or “moment 1” calculated after applying the strict mask to the 150 pc resolution data cube for NGC 4303. The right panel shows the intensity-weighted mean velocity calculated combining the strict and broad moment 1 maps. The broad map is only used where there is no strictly masked measurement, the integrated intensity exceeds an S/N of 2, and the velocity field lies within some tolerance ($\pm 30 \text{ km s}^{-1}$) of a prior estimate—in this case, a lower-resolution, strictly masked velocity field. This “moment 1 with a prior” significantly expands the coverage of the velocity field while still yielding coherent structure.

Figure 23. Products showing line widths. The figure shows two measures of line width for NGC 4303 calculated as part of our product creation. Left: rms velocity dispersion, also referred to as the “moment 2” map, calculated after applying the strict mask to the cube. Right: “Effective width” or “equivalent width” (Equation (5), Heyer et al. 2001) also calculated from the strictly masked cubes. The two maps show overall similar distributions, with disagreements arising in cases where the line profile is non-Gaussian and where noise affects the spectrum.

potential problems. For the initial internal PHANGS–ALMA data inspections, each report presented the following diagnostics:

1. Summary of the beam size, shape, and orientation and the astrometric grid. These parameters were extracted from the FITS header of the cube.
2. Channel maps showing the deconvolved data cube, user-defined clean mask (if any), data cube of residual emission, and the ratio of the data to the residuals.
3. A moment 0 map produced with no masking, i.e., generated by summing the full cube over the full imaged bandwidth.
4. Tables reporting the sum of the emission inside and outside the user-supplied clean mask and a mask identifying significant emission. This is not identical to the “strict mask” defined in Section 7, but it is constructed along similar lines.

5. Histograms of the pixel values in the full cube, and separate histograms for pixels inside and outside the user-supplied clean mask and mask identifying significant emission.
6. Integrated spectra, constructed by summing the full data cube, and separate spectra from summing the cube with the user-supplied clean mask and significant emission mask applied.
7. Two-dimensional histograms illustrating the distribution of pixel values within each channel for the full data cube, as well as versions for the residual image and the masked and unmasked regions of the cube.
8. Power spectra of emission calculated from the individual channel maps.

We found that this collection of plots allowed the reviewers to identify pathological data and to evaluate the performance of the deconvolution and masking.

In practice, the reports used for internal QA were generated by an IDL pipeline that ran independently of the PHANGS–ALMA data reduction pipeline. We subsequently developed a Python version of the QA report generation tool that can be run independently or integrated into the PHANGS pipeline.

8.2. Regression Tests against Previous Versions

The PHANGS–ALMA imaging and product creation pipelines have been iterated several times. These iterations included a major code revision and several substantive revisions of the imaging and product creation algorithms. For products created after the initial round of quality assurance, we could use the previous, already quality-assured imaging as a benchmark against which to compare the new products.

For these newer products, we automatically generated the QA reports as before, but we adopted a different, less time-consuming QA strategy. We created a suite of regression tests that benchmarked each new cube and map against the equivalent, quality-assured product created by a previous version of the pipeline. The statistics used for these regression tests included: beam shape, astrometric grid parameters, key statistics describing the distribution of pixel values, integrated flux, flux above fixed intensity and S/N thresholds, and outlier-resistant standard deviation estimates.

Using the regression tests, we checked whether each of the parameters extracted from the new and previous versions of the data were in good agreement. To define acceptable agreement, we imposed typical tolerance levels from 1% to 20%, depending on the parameter under consideration and our knowledge of the changes that had been implemented in the pipeline. Following these regression tests, we then focused our manual QA efforts—including detailed inspection of the QA reports—on cases where the regression tests indicated significant differences between the new and old data products. We found that this approach represented an acceptable compromise between rigorously testing the impact of each change to the PHANGS–ALMA pipeline on all PHANGS–ALMA data products and overwhelming our manual QA team.

8.3. End-to-end Test of the PHANGS Pipeline

We also tested the performance of the pipeline by applying it to simulated data. For this test, we created a series of simulated CO(2–1) measurement sets using CASA’s `simdata` task. We consider two source intensity distributions and assume the

same observing conditions across all simulated observations, but we vary the overall amplitude of the signal in each input model to create a suite of data sets with differing S/N. We then ran the simulated measurement sets through the staging, imaging, and postprocessing parts of the pipeline. Finally, we compared the output from the pipeline to the input model image to assess the performance of the pipeline.

Simulation setup: We simulated observations of CO(2–1) emission from two sources: (1) a modified, more distant version of NGC 1097, and (2) NGC 3059 with no modification to the distance. The two cases span the range of structure that we see in the real data set. The modified NGC 1097 has compact, bright structure in each channel, partially because it displays a very strong velocity gradient and hosts strong features in the form of a compact circumnuclear ring, a strong bar, and well-defined spiral arms. NGC 3059 shows more extended structure in individual channels and a more flocculent overall structure. Reflecting this, the real data for the two targets show different results when comparing the 12 m+7 m and 7 m only results in Appendix C. In NGC 1097, we find almost no discrepancy between the two cases, while NGC 3059 shows much higher flux in the cleaned 12 m+7 m compared to 7 m only data. We chose these targets in part to help us understand this effect (see more in Appendix C).

Specifically, we produced the model data cube following these steps:

1. We began with the “strictly masked” 12 m+7 m+TP CO(2–1) data cube for each galaxy. For NGC 1097, this cube combined two individual mosaics (“parts”), observed separately. NGC 3059 was observed in a single part.
2. We rotated each image to align the major axis of the cube with the decl. axis. We also resampled each cube to have channel width $\sim 0.6 \text{ km s}^{-1}$, using cubic interpolation in CASA’s `imregrid` to interpolate from the cube’s 2.54 km s^{-1} channels. Then, we converted to units of Jy pix^{-1} .
3. Only for NGC 1097, we adjusted the pixel scale of the image. In order to ensure that the fine-scale structure in our input model is sharper than our observed beam, we shrunk the pixel scale of the model cube by a factor of two, i.e., a factor of four in area. This effectively places the model image at two times the real distance to NGC 1097, moving it from 13.6 Mpc (Shaya et al. 2017; Anand et al. 2021) to 27.2 Mpc. During this step, we leave the intensity in Jy pix^{-1} unchanged. Thus, the initial NGC 1097 look-alike model can be thought of as a target with twice the distance and four times the luminosity of NGC 1097.
4. We do not apply any such rescaling to NGC 3059. This means that, in this target, we simulate observing structure that has already been convolved with the telescope beam. Because this case is intended to test the imaging performance for more extended sources, we do not view this as a problem.
5. In both targets, we added a continuum with 1/30 the peak intensity along each line of sight. This is significantly brighter than our typical continuum, but should be removed by our continuum subtraction.
5. We also created additional versions of each model by dividing the intensity in each pixel by 3, 10, 30, and 100. Because we use a fixed simulated observing time and

Table 6
End-to-end Imaging Tests

Flux ^a (Jy km s ⁻¹)	Model				
	1/1 Model ^b	1/3 Model	1/10 Model	1/30 Model	1/100 Model
NGC 1097 Look-alike					
Characteristic 7 m model S/N ^c	94	34	10	3.5	1.1
Model	6087 (100%)	2029 (100%)	609 (100%)	203 (100%)	61 (100%)
7 m pipeline clean	5931 (97%)	1901 (94%)	526 (86%)	149 (73%)	44 (69%)
12 m+7 m pipeline clean	6202 (102%)	2087 (103%)	610 (100%)	190 (94%)	60 (94%)
7 m+TP pipeline image	6111 (100%)	2045 (101%)	624 (102%)	217 (107%)	77 (120%)
12 m+7 m+TP pipeline clean	6112 (100%)	2050 (101%)	627 (103%)	220 (108%)	80 (125%)
NGC 3059 Look-alike					
Characteristic 7 m model S/N ^c	18	6.1	1.8	0.6	0.2
Model	924 (100%)	308 (100%)	92 (100%)	31 (100%)	9.2 (100%)
7 m pipeline clean	745 (81%)	219 (71%)	70 (76%)	42 (135%)	38 (413%)
12 m+7 m pipeline clean	941 (102%)	343 (111%)	153 (165%)	118 (381%)	111 (1200%)
7 m+TP pipeline image	934 (101%)	318 (103%)	103 (112%)	41 (132%)	20 (217%)
12 m+7 m+TP pipeline clean	937 (101%)	321 (104%)	106 (115%)	45 (145%)	23 (250%)

Notes. Summary of inferred flux (model input, clean flux, integrated flux) for our five simulated CO(2–1) measurement sets.

^a Integrated flux calculated after continuum subtraction in the model. For the 12 m+7 m and 7 m data, we report the total cleaned flux. For the feathered data, we report the flux in the final cube after feathering.

^b See text. The nominal 1/1 model is the version created as described in the text based on the strictly masked NGC 1097 or NGC 3059 imaging. The scaled versions reduce the intensity of all voxels by factors of 3, 10, 30, and 100.

^c Intensity-weighted intensity value of emission in the input model after convolution to the resolution of the 7 m, divided by the 1σ noise in that 7 m cube. This is a characteristic signal-to-noise value for the data, and gives an indication of the brightness of emission in the cube, though the detailed brightness distribution is complex and resolution-dependent.

fixed weather conditions, these different versions correspond to cases with the same structure but different S/N values. We report characteristic S/N values for each case in Table 6. In practice, the NGC 3059 cube scaled down by a factor of 100 yields no meaningful results, because the galaxy becomes too faint to be detected using our simulated observations.

We simulate interferometric observations of each model image using CASA’s `simobserve` task. In detail, we simulate observing for 6 hr using the ALMA Cycle 5 ACA and for 1.5 hr using the most compact Cycle 5 12 m configuration (i.e., C43-1). All simulations occurred around transit and included simulated thermal noise appropriate for 1 mm of precipitable water vapor. We allowed the simulator to place mosaic fields that would cover the target using the default spacing. For NGC 1097, the simulator placed 33 ACA 7 m pointings and 92 12 m pointings. For NGC 3059, the simulator placed 67 ACA 7 m pointings and 203 12 m pointings.⁵⁰

We also created corresponding simulated single-dish cubes. To do this, we convolved each model to the resolution of the single-dish data using the CASA task `imsmooth`. For this step, we did not include any continuum, in order to simulate the baseline subtraction during the single-dish processing (Section 5). Then, we added noise to each simulated single-dish cube. The noise that we added was first convolved to the resolution of the single-dish data and then scaled such that the rms amplitude of the simulated noise matched the measured 1σ noise level of the real NGC 1097 single-dish cube.

When these steps were finished, we had simulated $u-v$ data and single-dish cubes for NGC 1097 and NGC 3059 look-alikes with five S/N levels (see Table 6). These data resemble typical observations obtained within the PHANGS–ALMA survey. They allow us to assess the pipeline performance because they correspond to known input images.

Pipeline imaging: We configured the pipeline to process the simulated data in a manner that closely followed the real PHANGS–ALMA imaging. We staged the data, subtracted the continuum, regridded, and rebinned to a data set ready for imaging. Then, we imaged and cleaned each data set and applied the postprocessing steps described in Section 6. These steps included feathering with the simulated single-dish data.

Next, we convolved each model input image to the resolution of the pipeline-produced output image. Then, we reprojected the model to the astrometric and velocity grid of the pipeline-produced data. Thus, at the end of this process, we had 12 m+7 m, 12 m+7 m+TP, 7 m, and 7 m+TP pipeline-imaged, simulated images for both the NGC 1097 and NGC 3059 look-alike, each at five S/N levels.

Results: Figure 24 shows the peak intensity from the beam-matched, aligned, input models and the output from the pipeline. We plot results for the 12 m+7 m+TP imaging of our NGC 1097 look-alike scaled down by a factor of three and our brightest NGC 3059 look-alike. In both cases, the imaging shows excellent recovery of the detailed features and large-scale morphology of the input image.

Table 6 and Figures 25–27 show these results in more detail. Figure 25 shows the most basic result, the scaling between input model image (x -axis) and PHANGS pipeline output image (y -axis) for the 12 m+7 m+TP and 7 m+TP data for both source models and all scalings. Before the comparison, we match the resolution and astrometric grid of the model to that of

⁵⁰ This is slightly larger than the ALMA observatory limit of 150 pointings, but this should have no effect on the test.

Figure 24. Example results from end-to-end tests of the pipeline using simulated data. The left panels show peak intensity images constructed from model input for two of the cases used to test the performance of the pipeline. The right panels show the peak intensity maps constructed from the pipeline output image using the same data. To provide a fair visual comparison, the model data have been convolved to match the resolution and astrometric grid of the pipeline output data. All panels show results for the 12 m+7 m+TP data, i.e., imaging the combined 12 m+7 m data and then feathering with the simulate total power data. The top row shows results for our NGC 1097 look-alike with the original model scaled down a factor of 3 in intensity. The bottom row shows results for our NGC 3059 look-alike.

the pipeline output cube. We restrict the comparison to regions that are above $5\times$ the noise in the observed cube in the matched model image (dashed red line). Overall, the figure demonstrates excellent performance of the pipeline, with data from all models lying almost exactly along the line of equality. Note that, for the two lowest-brightness versions of the NGC 3059 look-alike, the S/N drops to such low levels that the deconvolution cleans only noise. These imaging cases fail because the simulated observations are not deep enough to see the galaxy.

Of course, the pipeline images do not perfectly match the input model, and Figures 26 and 27 explore the offset between the input and output in more detail. These figures plot the difference between the pipeline and model for individual voxels in the cubes. We show results for the brightest version of the NGC 1097 look-alike (Figure 26) and the brightest version of the NGC 3059 look-alike (Figure 27). In each panel, we plot individual (gray) and binned (black) data for the voxels

that account for 95% of the emission in the model image. That is, we construct a CDF of the flux as a function of intensity, and show results for the top 95%. The exact 95% value of the threshold is arbitrary; we only need some cut to select regions of interest where the model is positive. This figure also shows the median and scatter for the remaining pixels with a blue point. We use red lines and shading to indicate the 2σ level of the statistical, predominantly thermal noise in the data cube.

Both figures show the same good agreement seen in Figure 25. Differences between the model and the pipeline output remain small compared to the intensity value in the cube. We do find low-level systematic deviations, however. As expected, the thermal noise contributes to the scatter; this provides the main explanation for the width of the distributions in Figures 25–27. We observe a modest positive bias in our pipeline imaging of NGC 1097—but not NGC 3059. The sense of the bias is that the pipeline yields results that are biased high relative to the model by $\sim 2\%$. To see this, compare the binned

Figure 25. *Input model and pipeline output results.* Scatter plots showing the pipeline-imaged output (y-axis) as a function of the input model intensity (x-axis) for (top row) the NGC 1097 look-alike models and (bottom row) the NGC 3059 look-alike models and (left column) 12 m+7 m+TP and (right column) 7 m+TP results. In each case, the input model has been convolved and aligned to the resolution and astrometric grid of the output cube before comparison. Diagonal solid lines show equality. Vertical dashed lines show a typical 5σ noise value in the simulated image. We apply this 5σ value as a threshold in the model. Different colors show individual voxels for each scale version of the input model. Overall, the pipeline results closely match the input model. We investigate the scatter more in Figures 26 and 27.

results (black points) to the dashed black lines, which illustrate the case where the model has been multiplied (upward curve) or divided (downward curve) by 1.02. The binned data match the upward-curved line well, indicating that, to first order, a 2% positive multiplicative bias and thermal noise provide a good description of the pipeline’s output compared to the input model.

The bottom left panels in both figures show the 7 m only imaging results. For both input models, the 7 m only imaging shows a negative offset at a wide range of intensities. For NGC 1097, this offset appears mild, but for NGC 3059, the

pipeline image is $\sim 1\text{--}2\sigma$ lower than the input model over a wide range of intensities. This reflects the poorer imaging performance of the 7 m only data compared to the 12 m+7 m data discussed in Appendix C. As mentioned above, the difference between the 7 m imaging performance for the two targets is expected. For the real NGC 1097 imaging, the 12 m +7 m imaging and 7 m imaging agree well, and little flux is lost to spatial filtering. Meanwhile, for NGC 3059, the real data show significant differences between the 12 m+7 m and 7 m images.

Figure 26. Difference between pipeline output and model input for the NGC 1097 look-alike. The plots show the difference between pipeline output and model input values (y-axis) as a function of model intensity (x-axis) for imaging of our NGC 1097 look-alike model. The panels show results for four cases: imaging with (top left) the 12 m+7 m arrays together; (top right) the full 12 m+7 m+TP results, i.e., 12 m+7 m imaging with feathering; (bottom left) imaging with only the 7 m array; and (bottom right) the full 7 m+TP results, i.e., 7 m imaging with feathering. Gray dots show results for those individual voxels that contribute 95% of the total flux in the model after sorting by intensity. Black points and error bars indicate the median and 1σ scatter in the residual. The blue points show the median, 1σ (thick error bar), and 2σ (thin error bar) for lower-intensity voxels. The red lines show $\pm 2\sigma$ thermal noise in the image, and black lines show $\pm 2\%$ fractional scatter. This figure shows results of running our fiducial model through end-to-end imaging tests. The fiducial model resembles a brighter version of NGC 1097 at twice the actual distance to that galaxy. Overall, the agreement between the pipeline results and the model is excellent, but we do see evidence of a $\sim 2\%$ multiplicative bias such that the pipeline results appear high in all panels.

We also report the total fluxes in the model, the cleaned images, and the feathered images for each case in Table 6. For all of the feathered cases, the pipeline matches the model input within the uncertainty expected from directly summing the total power cube.⁵¹ This must be the case: the feathering operation will fix the integral of the cube to match the input simulated total power data, which in turn is simply the model galaxy plus noise.

Table 6 also gives insight into the performance of the deconvolution. For both the NGC 1097 and NGC 3059 look-alikes, the 12 m+7 m imaging recovers results within 2%–3%

% of the model flux, even before feathering. For lower-brightness versions of the NGC 1097 look-alike, the 12 m+7 m imaging continues to reconstruct almost all of the emission. For the NGC 3059 look-alike, as mentioned above, the imaging fails for the two faintest cases, which have typical 7 m S/N levels of ~ 0.6 and 0.2 .

The situation with the 7 m data is more mixed. In the NGC 1097 look-alike, the deconvolution recovers 94%, 86%, 73%, and 72% of the model emission for the models scaled down by factors of 3, 10, 30, and 100. For the NGC 3059 look-alike, the situation is even worse, with the 7 m only imaging recovering only 81% and 71% of the total flux in the two brightest cases. In short, for low S/N, ACA-only data sets, the PHANGS pipeline struggles to achieve a full deconvolution of

⁵¹ Note that, for expediency, we used the same noise realization for all TP simulations, so they are all biased high relative to the model by roughly the same $+20 \text{ Jy km s}^{-1}$.

Figure 27. Difference between pipeline output and model input for the NGC 3059 look-alike. As Figure 27, but for our NGC 3059 look-alike, reflecting a fainter galaxy with more extended source structure. This case does not show any notable bias with the 12 m+7 m data, but the 7 m only imaging shows a bias low. For the 7 m+TP imaging, a mild bias plus thermal noise represents a good model for the small offset of the pipeline results from the model. For the 7 m only data, the imaging struggles to recover the model. See Appendix C for more discussion.

the 7 m only data. In our end-to-end tests, the PHANGS pipeline ACA-only imaging misses 20%–30% of the flux in the worst cases. We explore this issue with both simulations and real data in Appendix C.

8.4. Comments on Quality Assurance Results

Both the automated regression and the manual quality assurance tests played an important role in refining the PHANGS pipeline algorithms and catching several important bugs. In the final round of imaging using the latest pipeline, we still found a few cases where the imaging with the default parameters diverged or declared convergence too early (~ 4 out of ~ 250 total cases) in one or more planes. In these cases, we adjusted the pipeline parameters and reran the imaging. Usually, adjusting the convergence criteria or the primary beam cutoff improved the situation.

9. Summary

We have presented the PHANGS–ALMA data processing pipeline, explaining the key steps in the processing and our motivation for many of the key decisions. We do not review these here, but do highlight a few points that may be of general interest to those working on similar problems:

1. We note that issues related to regridding and interpolation can lead to patterns in the rms noise amplitude along the spectral dimension. The specific issue that we highlight is related to CASA, and we hope it will be addressable in future releases. However, the concern that data processing affects the spectral noise pattern and line spread function is general.
2. We present a robust two-stage approach to deconvolving spectral line observations. This approach employs a multiscale deconvolution down to an S/N threshold of around four. It then creates a mask based on bright signal

and uses a classic Högbom (1974) deconvolution approach to clean the bright emission “into the noise.” We have found this to run robustly and yield good results on a wide variety of nearby galaxy data.

3. We adopt a two-track approach to masking of spectral line data cubes. We create a “strict” mask with high confidence and a low false-positive rate, but potentially low completeness. This is appropriate for calculations that perform poorly in the presence of noise. We also create a high-completeness but noisier “broad” mask that will include many false positives but also encompass almost all emission in the cube. Leaving aside the specifics of their creation, we suggest that this two-track approach to masking is a good general approach.
4. We have found that a “matched-line width filtered peak temperature map” does an outstanding job of highlighting detailed structure in line data cubes. This is simply a conventional peak temperature (sometimes referred to as “moment 8”) map constructed from a cube that has been convolved spectrally with a matched-line width filter.
5. We have compared the imaging results for different arrays and vetted the performance of our pipeline using simulated observations with known input. These tests show that, after the inclusion of total power data, the pipeline does an excellent job of recovering known input. They also show that the 12 m+7 m imaging performs significantly better than 7 m only imaging in many cases, even after matching the resolution of the output images. The differences, which are explored in detail in the appendix, are a function of S/N and source structure.

The PHANGS–ALMA pipeline has so far been applied successfully to roughly 1000 individual measurement sets, and is publicly available on [github](https://github.com/akleroy/phangs_imaging_scripts).⁵²

We gratefully acknowledge a prompt and constructive review by the anonymous referee during a difficult time. This work also benefited immensely from helpful discussions with Crystal Brogan, Jeffrey Mangum, and the North American ALMA Science Center and European Southern Observatory staff, including Dirk Petry. The computing and software infrastructure used to process PHANGS–ALMA and develop this pipeline at OSU was built and supported by David Will, who for years was the most supportive and welcoming person in a supportive and welcoming department. He will be sorely missed.

This work was carried out as part of the PHANGS collaboration. The work of A.K.L., J.S., and D.U. is partially supported by the National Science Foundation (NSF) under grants No. 1615105, 1615109, and 1653300, as well as by the National Aeronautics and Space Administration (NASA) under ADAP grants NNX16AF48G and NNX17AF39G. E.R. acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), funding reference number RGPIN-2017-03987, and computational support from Compute Canada. D.L., T.S., E.S., C.M.F., K.S., and T.G.W. acknowledge funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 694343). C.H., A.H., and J.P. acknowledge support by the Programme National “Physique et Chimie du Milieu

Interstellaire” (PCMI) of CNRS/INSU with INC/INP co-funded by CEA and CNES. A.H. acknowledges support by the Programme National Cosmologie et Galaxies (PNCG) of CNRS/INSU with INP and IN2P3, co-funded by CEA and CNES. A.U. and A.G.-R. acknowledge support from the Spanish funding grants AYA2016-79006-P (MINECO/FEDER) and PID2019-108765GB-I00 (MICINN). A.U. acknowledges support from the Spanish funding grant PGC2018-094671-B-I00 (MCIU/AEI/FEDER). C.M.F. acknowledges support from the NSF under Award No. 1903946. M.C. and J.M.D.K. gratefully acknowledge funding from the German Research Foundation (DFG) through an Emmy Noether Research Group (grant No. KR4801/1-1). M. C., J.M.D.K., and J.J.K. gratefully acknowledge funding from the DFG Sachbeihilfe (grant No. KR4801/2-1). J.M.D.K. gratefully acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program via the ERC Starting Grant MUSTANG (grant agreement No. 714907). F.B., A.T.B., I.B., J.d.B., and J.P. acknowledge funding from the European Union’s Horizon 2020 research and innovation program (grant agreement No. 726384/EMPIRE). R.S.K. and S.C.O.G. acknowledge financial support from the DFG via the collaborative research center (SFB 881, Project-ID 138713538) “The Milky Way System” (subprojects A1, B1, B2, and B8). They also acknowledge subsidies in the framework of Germany’s Excellence Strategy (grant EXC-2181/1-390900948) and funding from the ERC via the ERC Synergy Grant *ECOGAL* (grant 855130). K.K. and F.S. gratefully acknowledge funding from the DFG in the form of an Emmy Noether Research Group (grant No. KR4598/2-1). E.W. acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 138713538—SFB 881 (“The Milky Way System,” subproject P2). C.E. acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG) Sachbeihilfe, grant No. BI1546/3-1. A.S. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-1903834.

This paper makes use of the following ALMA data, which have been processed as part of the PHANGS–ALMA CO(2–1) survey: ADS/JAO.ALMA#2012.1.00650.S, ADS/JAO.ALMA#2013.1.00803.S, ADS/JAO.ALMA#2013.1.01161.S, ADS/JAO.ALMA#2015.1.00121.S, ADS/JAO.ALMA#2015.1.00782.S, ADS/JAO.ALMA#2015.1.00925.S, ADS/JAO.ALMA#2015.1.00956.S, ADS/JAO.ALMA#2016.1.00386.S, ADS/JAO.ALMA#2017.1.00392.S, ADS/JAO.ALMA#2017.1.00766.S, ADS/JAO.ALMA#2017.1.00886.L, ADS/JAO.ALMA#2018.1.01321.S, ADS/JAO.ALMA#2018.1.01651.S, ADS/JAO.ALMA#2018.A.00062.S. ALMA is a partnership of ESO (representing its member states), NSF (USA), and NINS (Japan), together with NRC (Canada), NSC and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO, and NAOJ. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

Software: ALMA Calibration Pipeline (L. Davis et al. 2021, in preparation), CASA (McMullin et al. 2007), numpy (Oliphant 2006), scipy (Virtanen et al. 2020), astropy (Astropy Collaboration et al. 2013, 2018) IDL Astronomy User’s Library (Landsman et al. 1993), cprops

⁵² https://github.com/akleroy/phangs_imaging_scripts

(Rosolowsky & Leroy 2006), GILDAS Pety et al. (2005), R R Core Team (2015), PHANGS–ALMA Total Power Pipeline (Herrera et al. 2020), spectral-cube, radio-beam.

Appendix A Contributions

The processing of the PHANGS–ALMA data and creation of the pipeline was a team effort, with major contributions from many people and input from the entire team. This paper also reflects major direct and indirect contributions from many people. We summarize some of the key contributions here.

The PHANGS ALMA Data Reduction (ADR) Group: The group has been led by J. Pety since the beginning of the PHANGS collaboration, and met weekly for most of the time since 2016. Key contributors to tests, discussions, and development over the course of the project include: M. Chevance, C. Faesi, C. Herrera, A. Hughes, A. Hygate, D. Liu, A. Leroy, T. Saito, E. Rosolowsky, E. Schinnerer, K. Sliwa, A. Schruba, and A. Usero.

Interferometric and Postprocessing Pipeline: The code was mostly developed by A. Leroy, D. Liu, E. Rosolowsky, and T. Saito, with code review at several stages by A. Schruba, and key input from J. Pety, E. Schinnerer, A. Schruba, and A. Usero. Additional tests related to many aspects of data processing were carried out by D. Liu and T. Saito. E. Koch and C. Wilson offered important input on algorithms. Tests of short-spacing correction algorithms were led by K. Sliwa during the pilot programs and then T. Saito during the Large Program and beyond. T. Saito led the research described in Appendix D of this paper. A. Leroy and T. Saito led the work described in Appendix C with major input from A. Hughes, J. Pety, E. Rosolowsky, and E. Schinnerer. The pipeline was deployed for PHANGS–ALMA by A. Leroy and T. Saito.

Total Power Pipeline: C. Herrera developed most of the total power pipeline, with major input from J. Pety and A. Usero, and code review by E. Rosolowsky. K. Sliwa played a key role in prototyping approaches to the total power processing using the pilot data. A. Usero developed and deployed the telluric ozone correction algorithm described in Section 5, and led investigation and communication of this issue in close collaboration with C. Faesi, C. Herrera, and J. Pety. A. Usero also led investigation of the calibration stability and gain in the total power data, described in Section 3.2 and Appendix B. A. Weiss, J. Pardo, and C. de Breuck also provided important input on this topic. A. Usero led investigation of the flux stability of the total power observations. The total power pipeline was deployed on PHANGS–ALMA by C. Faesi, C. Herrera, and A. Usero.

Quality Assurance (Cubes): A. Hughes developed the IDL version of the quality assurance software for cubes described in Section 8. D. Liu wrote the Python version. A. Hughes led regression testing and, with J. Pety, coordinated cube quality assurance efforts at several stages. A. Leroy and T. Saito developed and deployed the end-to-end tests described in Section 8.3. Many members of the team contributed careful review of data products, including: I. Beslic, M. Chevance, J. den Brok, C. Eibensteiner, C. Faesi, A. García-Rodríguez, C. Herrera, A. Hygate, M. Jimenez Donaire, J. Kim, A. Leroy, D. Liu, J. Pety, J. Puschnig, M. Querejeta, E. Rosolowsky, T. Saito, E. Schinnerer, A. Schruba, A. Sardone, J. Sun, A. Usero, D. Utomo, and T. Williams.

Quality Assurance (Visibility Data): T. Saito developed the $u-v$ quality assurance procedures and software described in

Section 3.2. T. Saito and C. Herrera carried out most of the inspection using these tools. D. Liu carried out the analysis of the flux calibration scale described in Section 3.2.

Infrastructure: E. Rosolowsky created and maintained the PHANGS server and shared archive, which was crucial to distributing the work and results. D. Will created and maintained the computing and software environments used for building and executing the pipeline at OSU.

Observatory and Community Support: The Joint ALMA Observatory and North American ALMA Science Center offered extensive support. They were responsive and flexible regarding reobservations of total power data affected by the telluric ozone feature (Section 5), and responsive and helpful when issues related to imaging in CASA arose early in the project. We specifically acknowledge helpful communication with A. Remijan regarding imaging, C. Brogan regarding the noise behavior in Band 6, and J. Mangum regarding several aspects of data processing. More broadly, this work builds on the hard work of the CASA team, the ALMA pipeline team, and the ALMA observatory effort to provide excellent quality assurance. We acknowledge the hard work of the developers, the scientists who support and guide the effort, and the data analysts. Similarly, we build on the large work by the *astropy* and broader scientific Python community, and also acknowledge the astronomical IDL community, which laid the foundation for much of this work.

Appendix B Internal Stability of the Calibration for the PHANGS– ALMA Total Power Data

The total power observations set the overall flux in our final data cubes. Our processing assumes that both the total power and interferometric data are correctly calibrated (Section 6.4 and Appendix D). The observatory calibration scheme anchors the amplitude calibration of the total power data to interferometric observations with the 7 m array, and so to the ALMA calibrator database. In principle, this should yield stable, high-quality calibration.

To check this, we assessed the internal consistency of the PHANGS–ALMA total power observations. Over the course of our ALMA Large Program, we observed the same targets repeatedly, and the individual observations are already deep enough that they detect most galaxies at high significance. This allows us to compare how the line brightness toward the same target varies when measured on different days. The magnitude of these variations gives us an upper limit to the stability of the ALMA total power flux calibration.

For this test, we selected a subset of six galaxies. Using the procedures described in Herrera et al. (2020) and Section 5, we generated a CO(2–1) line cube for every “execution block” (EB), each of which corresponds to an individual ~ 1 hr long observation. This resulted in 31 independent data cubes, with $N = 2-7$ cubes per galaxy. Each target galaxy was observed on 2–3 different dates, with the spread between observations spanning from 4 days to 11 months. Typically, two consecutive EBs were observed on any given day.

For each galaxy, we consider every possible pair of EBs. For each pair, we fitted a linear function with no intercept to the scatter plot of intensities measured at the same voxel position in the two data cubes. For the fit, we used a total-least-squares linear regression scheme, taking into account the noise level in both cubes. Because this exercise requires comparing detected

Figure 28. Assessment of the internal calibration of ALMA’s total power observations. Each histogram illustrates one of the four regression tests described in the text that we run on 74×2 EB pairs. The top and bottom panels show the two tests on line cubes and moment 0 maps, respectively. The blue filled histograms correspond to tests where all valid data are taken. The black histograms correspond to tests including only data where the nominal S/N in the ratio between the two measurements is >10 . The orange dotted-dashed curve is a reference Gaussian distribution with 3% rms (0.013 dex).

emission between the two EBs, we only consider velocities within the galaxy’s velocity range.

As a crosscheck, we repeated the exercise using the integrated intensity (moment 0) maps obtained by integrating over the galaxy’s velocity range. We also carried out fits both using all data and restricting to data with $S/N > 10$ in the ratio Y/X , where Y and X refer to the intensities or integrated intensities in the two data sets. We also repeated each comparison considering $X|Y$ and $Y|X$, i.e., swapping which data set we treat as the reference variable, in order to prevent any biases from the way EBs are processed. Thus, we ended up with four different slope measurements: maps and cubes, each with and without an S/N cut. In total, we compared 74×2 EB pairs, and also considered each with $Y|X$ and $X|Y$.

Figure 28 shows the distribution of fit slopes from each of our four tests. We expect a slope of unity ($\log_{10}$ fit slope = 0) if the calibration is identical between the two EBs. Departures from this capture the variations in the relative calibration between the two EBs. We report results in log scale because it is natural to expect calibration uncertainties to be multiplicative.

The rms of the distribution of fit slopes shown in Figure 28 is 0.010–0.017 dex ($\approx 2\%$ – 4%), regardless of how we calculate it: directly or by fitting the histograms with a Gaussian function. Although the histograms hint at some low-level wings, almost all the slope measurements deviate by less than $\sim 7\%$ (0.03 dex) from unity. We adopt a 3% rms scatter on the EB-to-EB scaling as a reasonable description of our results. We did not find any convincing indication that the slopes depend on the difference between the observing dates of the EB pairs. This suggests that the calibration uncertainties affecting different EBs are mutually uncorrelated.

We interpret this scatter as indicative of the stability of the ALMA total power calibration. It will not reflect underlying

uncertainties in the ALMA calibrator database, but most other uncertainties should be captured by this test. Assuming that calibration uncertainties affecting different EBs are uncorrelated, the uncertainty in individual EBs is about $\sqrt{2}$ times lower than the measured 0.03 dex $\approx 3\%$ difference between two independent EBs.

The final total power cube for each galaxy typically results from averaging several EBs. If the calibration uncertainty is uncorrelated, this would reduce the uncertainty from 3% by an additional $\sqrt{N}$ factor, where N is the number of EBs. Thus, we conclude the internal calibration of the PHANGS–ALMA data is robust at the $\sim 1\%$ level. Given the link to the interferometric calibration scheme, this experiment also bolsters our confidence in adopting the observatory-provided calibration without rescaling when combining the total power and interferometric data.

A note on absolute calibration for PHANGS–ALMA “version 4”: As part of our quality assurance, we examined the stability of the total power antenna gain. This is the observatory-provided number used to translate from the “chopper wheel”-based Kelvin scale to an absolute flux scale. It is expressed in units of Jansky-per-Kelvin, or Jy K^{-1} . We found that the observatory-provided Jy K^{-1} of individual observations varied systematically by $\sim 7\%$ based on delivery date. Consultation with the ALMA observatory revealed that this reflects delivery of some incorrect gain values during the time period 2017–2018. Surface improvements to the total power telescopes improved the gain, but there was some lag in reflecting these improvements within the delivered products. The correction for this effect is straightforward, but requires reprocessing the data from the original execution block stage, and so will not be reflected in PHANGS–ALMA’s initial public delivery, “version 4.” Taking into account the averaging of multiple blocks, we estimate that this effect implies a 2%–5% bias high for the overall flux scale of the data set for data delivered during 2018. The issue was not severe enough to disrupt the internal stability tests described here, and we expect it to be addressed in future releases.

Appendix C

Relative Performance of 7 m and 12 m+7 m Imaging

We image most of our targets using both the 7 m only and the combined 12 m+7 m data. In Section 8.3, we do the same with simulated galaxies in order to test the performance of the pipeline. These tests consistently show that the 7 m only imaging tends to deconvolve less flux than the 12 m+7 m imaging. The effect appears strongest for extended sources and at modest S/N.

We demonstrate this effect for both the real PHANGS–ALMA data and the simulation in Figures 29 and 30. The left panel in Figure 29 shows the ratio of cleaned flux in the 12 m+7 m imaging to that in the 7 m only imaging. That is, the y-axis shows the ratio of the summed fluxes of the model image. The x-axis shows a quantity that traces the typical S/N of reconstructed emission in the image. Specifically, we calculate the intensity-weighted mean intensity in the 7 m model via:

$$\langle I_\nu \rangle = \frac{\sum_{\text{model}} I_\nu^2}{\sum_{\text{model}} I_\nu}, \quad (\text{C1})$$

where I_ν is the intensity of a voxel in the model. Then, $\langle I_\nu \rangle$ is just the weighted average intensity in the model, and so captures the typical brightness in the image. In Figures 29 and 30, we plot $\langle I_\nu \rangle$ on the upper x-axis. The lower x-axis shows

Figure 29. *Deconvolved flux in different arrays.* (Left) The ratio of cleaned flux between the 12 m+7 m and 7 m only array images of the same galaxy (y-axis) as a function of the typical signal-to-noise ratio of deconvolved emission in the 7 m model image (x-axis; see text). (Right) The ratio of intensity at the peak of the 12 m+7 m model to intensity at the same location in the 7 m model, after convolving both to a matched $10''$ resolution. Again, we plot this as a function of the approximate typical signal-to-noise in the 7 m model image. Blue points in both panels show PHANGS–ALMA targets. Red points show results for the two sets of simulations described in Section 8.3, with the bright, compact NGC 1097–based galaxy appearing as a star and the more extended NGC 3059–based galaxy as a hexagon. The 12 m+7 m imaging recovers more flux than the 7 m only imaging, and the difference depends on the typical signal-to-noise of the emission. Both the integrated flux and the peak flux show this effect, but the effect is much stronger with regards to the integrated flux.

$\langle I_\nu \rangle / \sigma$, where $\sigma = 0.031 \text{ Jy beam}^{-1}$ is a typical noise level across our full 7 m imaging data set. In both figures, we use blue points to plot each galaxy part imaged with both 12 m+7 m and with 7 m only. The simulations described in Section 8.3 appear as red points, with different shapes reflecting our two model distributions at different S/N levels.

Imaging the 7 m only data deconvolves less flux than imaging the 12 m+7 m data: The left panel of Figure 29 shows that, in both the simulations and the real data, the 12 m+7 m imaging consistently deconvolves more flux than the 7 m only imaging. The effect shows a clear anticorrelation with the typical S/N in the data, such that bright targets show a much better match between the two deconvolved images than fainter targets. For low-brightness sources, the 7 m model can contain as little as $\sim 60\%$ of the 12 m+7 m flux. Across the whole sample, the 7 m only imaging deconvolves a median 73% of the flux deconvolved by the 12 m+7 m imaging.

The red points in Figure 29 show that this effect occurs in the simulated data, too. The deconvolved fluxes in Table 6 also highlight this result. The simulations also show significant spread at fixed S/N, demonstrating that source structure plays a large role in deconvolution. The NGC 3059 look-alike galaxy has extended structure within individual channels and shows more severe discrepancies at a given S/N than the compact NGC 1097 look-alike. Similarly, the large spread in the real data at any given S/N level likely reflects differences in the source structure within individual channels. Consistent with this interpretation, the model for our compact simulation, NGC 1097, shows some of the best agreement between the 12 m+7 m and 7 m only data across the whole sample.

The right panel in Figure 29 shows that the effect is present, though with smaller magnitude, even at the peak of the map. Here, we plot the ratio of peak intensities in the two models after convolving both to $10''$ resolution. On average, the 7 m only imaging achieves a peak intensity $\sim 93\%$ of that found

using the 12 m+7 m data at matched scales. Again, we observe an anticorrelation of the effect with S/N, but with significant source-to-source scatter.

The 7 m only imaging shows these shortcomings despite the fact that we clean the 7 m data to the 1σ level during the single-scale clean (Section 4), and none of our by-hand attempts to improve the cleaning yielded systematically better results for the 7 m only data. Although we do require nearby 4σ emission to conduct the single-scale clean, our by-eye assessment of the residuals does not reveal any systematic isolated $<4\sigma$ emission that could explain this behavior. Similarly, the 7 m images produced by our pipeline appear to compare favorably to the observatory-delivered products. In short, we have no reason to think that a simple algorithmic fix can address the issue, despite exploring several options. On the other hand, the 12 m+7 m data show excellent match to the input model in simulations and do not exceed the overall flux constraints set by the total power data. These 12 m+7 m images do, in fact, appear to represent our best images, and the combined arrays do a better job at flux recovery than the 7 m data alone.

Our best explanation for the issue is that the 12 m+7 m data have significantly better sensitivity on the relevant spatial scales to reconstruct emission from galaxies. Even at short $u-v$ separation distances, the 12 m only baselines add significant sensitivity and lead to a synthesized beam with fewer strong sidelobes. The 7 m data, on the other hand, have less sensitivity on scales matched to the emission and poorer $u-v$ coverage than the 12 m+7 m data. As a result, our deconvolution recovers less flux in the 7 m only image than the 12 m+7 m image.

While this qualitative explanation seems reasonable, we were surprised by the magnitude of the effect in the real data. Our current understanding is that, for realistic structures in nearby galaxies, the nonlinear nature of the deconvolution procedure interacts in a destructive way with the limited $u-v$ coverage and sensitivity of the 7 m array. This agrees qualitatively with

Figure 30. *Sum and peak flux in different arrays with and without feathering.* Similar to Figure 29, but now showing (top row) the ratio of the total summed flux in the 12 m+7 m image to that in the 7 m only image (top left) without feathering and (top right) with feathering. The bottom row shows the ratio of 12 m+7 m to 7 m intensities at the peak of the 12 m+7 m image after matching the resolution of the two images at $10''$. The bottom left image shows the case before feathering. The bottom right image shows the case after feathering. Again, the blue points show results for PHANGS–ALMA targets and red points show simulation results (Section 8.3). The 12 m+7 m imaging recovers more flux than the 7 m only imaging, with the two becoming more consistent as the signal-to-noise increases. After feathering, the two images are much more consistent for targets with reasonably high signal-to-noise, but still show some divergence in very low signal-to-noise cases. This effect appears much stronger in the integrated emission, but is present at a lower level in the peak intensity of the image. Both the simulations and real data show the effect, though it appears stronger in the real data due to the idealized nature of the simulations.

previous investigations on similar topics, which show a strong nonlinearity in interferometric image reconstruction of complex sources when dealing with limited coverage and sensitivity (in particular, see Helfer et al. (2002)). The fact that the idealized simulations show results qualitatively similar to the data underscores that this is not an effect driven solely by calibration issues, limited knowledge of the ALMA primary beams, or some similar issue that would only affect the observations. These issues may still play important secondary roles, however. In particular, we expect that the combined impact of phase and amplitude noise might also lead to nonlinearities in the

sensitivity of the interferometer to any extended structures at low S/N (e.g., see Lay et al. (1994), for a discussion of the nature of amplitude noise). More careful analysis of low S/N simulations with phase noise might help this situation.

In the real data, the images show the same effect: The top left panel of Figure 30 shows the ratio of total flux in the 12 m+7 m image to that in the 7 m only image. The lower left panel shows the ratio of peak intensities between these two images at matched $10''$. In other words, the left column of Figure 30 shows the same results as Figure 29, but for the real data, with the residuals added back into the deconvolved model.

Including the residual emission improves the situation by a small amount for the real data. However, a significant overall offset between the total flux in the 12 m+7 m and 7 m images remains visible in the top left panel of Figure 30. The simulations show a much larger improvement when the residuals are included. We understand this to reflect that the residuals in the simulations are idealized relative to those in the real data. We expect this because our simulations neglect phase noise. The simulations also achieve better rotation synthesis, and so better $u-v$ coverage, than the real data because the simulations observe a long continuous 7 m block around transit while the real observations observe short blocks at random times. Still, the simulations continue to show a significant effect in the final images. The lower panel shows that effect also appears for the peaks of the image in the real data.

Feathering largely corrects the issue, but not at the lowest S/N levels: In PHANGS–ALMA, we always attempt to include total power observations. These serve to anchor the total flux in the images and also to provide short-spacing information in the final image. The right column of Figure 30 shows the ratios of flux (top) and peak intensity (bottom) between the 12 m+7 m and 7 m data after feathering.

The figures show that, in cases with mean intensity levels $\gtrsim 4\sigma$, the flux discrepancy almost vanishes after feathering. In the lower-significance cases, we expect that the overall S/N in the cube is so low that statistical uncertainties in, e.g., masking to calculate the total flux, may drive some of the visible scatter. The bottom right panel shows that the agreement in the peak intensity also becomes much better with feathering, implying that the short-spacing correction helps locally, not only for the global flux. Overall, feathering reduces discrepancies in the integrated flux to $\lesssim 5\%$.

Synthesis: These results clearly demonstrate that the deconvolution of 7 m only data suffers from significant shortcomings and that these persist into the final images for the real 7 m only data. Because the discrepancies are largely resolved by feathering, we expect our final data products, even those involving only the 7 m array, to be largely correct. However, 7 m only images of nearby galaxies should be viewed with caution. Even the feathered data seem likely to harbor second-order fidelity issues, though a detailed investigation beyond what we present in Section 8.3 will have to wait for future work.

This analysis highlights a somewhat unexpected point. It would have been easy to attribute the flux missed from the 7 m deconvolution to “spatial filtering” that can only be addressed by including short-spacing data. This does not appear to be the case. Instead, in galaxies with clumpy structure and strong velocity gradients, including sensitive 12 m data significantly improves our flux recovery.

Appendix D Testing Short-spacing Correction Methods

Correcting interferometric data for missing short and zero spacings is a key part of reconstructing the true intensity distribution on the sky. Several SSC methods have been proposed, and there is not yet a clear consensus on the best approach. In this appendix, we test the suitability of three of these methods for PHANGS–ALMA. Our goal here is not a thorough assessment of each method, which would require a large amount of research; e.g., see ALMA memos 398 and 488 (Pety et al. 2001b; Tsutsumi et al. 2004). Instead, we test how

three popular methods work in a test case constructed to suffer from extreme spatial filtering.

To conduct this test, we use the CASA task `simobserve` to simulate interferometric and total power observations of artificial intensity distributions. We then image and combine these observations using different SSC techniques. We compare the recovered image to the known input image in order to evaluate the performance of each SSC technique.

We conduct these tests using a kind of “worst case” scenario for spatial filtering in a data set like PHANGS–ALMA. For our model input images, we adopt peak intensity (see Section 7) maps derived from the real PHANGS–ALMA data. We apply some clipping to isolate significant emission, but remove all velocity structure from the map. This leads to a model with widespread positive emission across each entire mosaic (e.g., as seen in the top left panels in Figures 31 and 32). This does not represent a truly realistic simulation of a galaxy. Our real targets show clumpy, sharp structure in individual velocity channels, while these models often show extended, relatively smooth structure on the scales accessed by the 7 m array used to carry out these tests. However, these models should present a case that can serve as a useful test of short-spacing correction algorithms. In that sense, this calculation complements the more realistic simulations in Section 8.3, in which the interferometric imaging recovers a larger fraction of the flux seen by the total power data.

Simulations: The simulations were set up as follows.

1. *Sample:* We run simulations considering 64 PHANGS–ALMA targets that had full 12 m+7 m+TP CO(2–1) imaging available and clear CO detections as of Fall 2019.
2. *Sky Model:* For each galaxy, we use the 12 m+7 m+TP peak CO(2–1) intensity map (Section 7) as the input sky model for `simobserve`. Before inputting them to the simulation, we clipped these images at a threshold corresponding to three times the rms noise in the map. By eye, this did a reasonable job of mostly including real structure emission associated with the galaxy. Pixels with values below this threshold had their values set to zero. After clipping, we convert the units of the maps to Jy beam^{-1} , appropriate for use with `simobserve`. We keep the sky coordinates the same as the true galaxy. For convenience, we set the source velocity to 0 km s^{-1} . This choice should have no impact on the imaging.
3. *Interferometric Simulations:* We use CASA 5.4.0 for the simulations, employing the task `simobserve` to construct the simulated measurement set and `tclean` for the subsequent imaging. We define our own hexagonally spaced mosaic grid to cover each input image. The spacing between neighboring pointings is set to $24''$, i.e., half-beam sampling for the 7 m array. For this exercise, galaxies observe with multiple independent mosaics in the real data set, e.g., NGC 2903 (Figure 11), were treated as a single image and not simulated in separate parts.

We simulate observations with the Cycle 5 ACA 7 m array configuration (`aca.cycle5.cfg`⁵³), reflecting the fact that most PHANGS–ALMA 7 m observations were obtained during Cycle 5 in the context of our Large Program.

The integration time of each $u-v$ data point is set to

⁵³ <https://almascience.eso.org/tools/casa-simulator>

10 s. The total observing time is set to 4 hr, which represents a reasonable match to the typical PHANGS–ALMA 7 m observing time per target. The source transits at the midpoint of the simulated observations. The observations thus happen at the highest possible elevation for each source.

We did not add thermal or phase noise. Instead, we concentrate on evaluation of SSC methods in the case of ideal observations. Aside from simply creating scatter in the measurements, we expect that the main effect of adding thermal and phase noise would be to add uncertainty to the deconvolution procedure. Thus, we prefer to focus on only the short-spacing correction here.

4. *Interferometric imaging*: We image the simulated visibility data using `tclean`. As in the PHANGS–ALMA pipeline (see Section 4), we start with a multiscale deconvolution using `tclean`, setting the `scale` parameter set to `[0, 2, 4]` pixels, with the pixel scale set to $1''$. We clean until the peak residual is ≤ 4 times the noise determined from the input image before clipping. We then continue with a single-scale (“Högbom”) `tclean`. This proceeds until the peak of the residuals ≤ 1 times the rms noise from the original image.

The imaging adopts Briggs weighting with `robust = 0.5`. We use `cyclefactor = 4`, `cell = 1''`, `cycleniter = 100`, and `gain = 0.2` (resp. 0.1) for multiscale clean (resp. single-scale clean). We do use a clean mask, which we create by convolving the input image to the synthesized beam size and detecting the significant pixels in this smoothed image.

5. *Total power data for simulations*: We simulate ideal single-dish observations by convolving the input sky image with the beam of the total power telescopes, $28''.6$ at 230.5 GHz.

Short-spacing Correction: There are currently at least three popular methods for short-spacing correction for ALMA. CASA’s `feather` routine represents the path recommended by many observatory guides and documentation. Alternatively, the `tp2vis` method offers the most advanced current implementation of joint deconvolution. A third path uses the input clean model to incorporate information on extended emission. Variations exist on each of these techniques, but broadly speaking, they span the range of current approaches in wide use. Before describing our results, we briefly describe each.

1. *Joint Deconvolution* (`tp2vis`): Koda et al. (2019) present and give a full description of the `tp2vis`. This method converts a total power map into visibilities via a “simple” deconvolution of the total power beam from the data in the Fourier plane. The data are then multiplied by the primary beam of the interferometric dish in the image plane—in this case, the 7 m antenna. Finally, the Fourier plane is fully sampled to produce a visibility table that reflects the total power data. The weight density of the total power visibilities is adjusted to match that of the interferometric visibilities. The total power visibilities are then merged with the interferometric ones using the CASA task `concat`. We use the same imaging scheme described above to image the combined interferometric and total power data set produced by `tp2vis`.
2. *Model-assisted CLEAN* (`tpmodel`): This method utilizes a user-defined image as an initial CLEAN model. This

has been used in various forms for several years (see, e.g., Dirienzo et al. 2015).

In our implementation, we pass the simulated total power image to `tclean` via the `startmodel` parameter. This initializes the CLEAN model, i.e., the deconvolved image, to the total power image. This input model is Fourier transformed and subtracted from the interferometric visibilities before the imaging proceeds. The imaging proceeds as normal, modifying the initial model until it achieves a good match to the interferometric data. Thus, this procedure essentially gives priority to the interferometer data and uses the total power as an additional guess to fill in missing information.

In practice, we first convert the units of the input total power image from Jy beam^{-1} to Jy pixel^{-1} , because CASA tracks the CLEAN model in these units.

3. *Feathering* (`feather`): In feathering, the total power and interferometric images are combined in the Fourier domain after imaging and deconvolution (e.g., Bajaja & van Albada 1979; Cotton 2017). The main difference between `feather` and `tp2vis` is that, with `feather` data, combination happens after deconvolution, so `feather` does not represent “joint” deconvolution. The advantage of this approach is its simplicity and robustness. This combination approach is adopted in much of the CASA documentation and has been commonly used in MIRIAD as well, where a version is implemented as `immerge`.

In practice, we run `feather` following the default CASA approach, matching the CASA guides. We set `sdfactor = 1.0`, `effdishdiam = -1.0`, and `lowpassfiltersd = False`. As in the main PHANGS–ALMA pipeline processing, we apply the primary beam correction before feathering. We found that this leads to better recovery of the source structure near the edge of the mosaicked field of view.

For each target, we use all three methods to create a short-spacing corrected image. The `tp2vis` method produces images with a slightly larger synthesized beam size than the other two methods. Therefore, we convolve all of the short-spacing corrected images to $10''$, in order to allow a direct comparison of the methods.

Evaluation of Results: Figures 31 and 32 show the results from our experiment for two galaxies: NGC 0628 and NGC 4303. The figures show the convolved input images, the output of imaging using only the 7 m data with no SSC, and the results for each of the three SSC approaches. All panels have the same beam and intensity scale, to allow direct comparison.

The top rows of Figures 31 and 32 show examples of the input model and output imaging, both convolved to $10''$ resolution, or slightly coarser than the typical resolution of the 7 m array at 230 GHz. In both targets, the 7 m only image shows extended negative sidelobes, or “negative bowls,” in place of extended positive emission in the input image. This illustrates how our choice to collapse the emission to a single plane yields an image that appears positive essentially everywhere. In both galaxies, significant, positive emission pervades the map at the resolution of the 7 m array. As a result, the interferometric images visible in the top middle panel show much stronger spatial filtering than we observe in our actual PHANGS–ALMA data. This highlights the fact that the total

Figure 31. Results of short-spacing correction tests for NGC 0628. The top left panel shows our model image convolved to the comparison resolution of $10''$. We construct the model image from a clipped version of the peak intensity map of the real combined 12 m+7 m+TP PHANGS–ALMA CO(2–1) observations for this target. Because we drop all velocity information for this exercise and focus on analysis of relatively low-resolution 7 m observations, the model shows positive emission essentially everywhere, with very smooth structure. The top middle panel shows the result of simulated observations using this input model and only the ACA 7 m array with no total power information. Significant missing flux can be seen in the image, indicating strong spatial filtering by the interferometer. The bottom row shows short-spacing corrected simulated observations. Each panel shows a different SSC method, from left to right: joint deconvolution using *tp2vis*, CLEANing using the simulated total power data as a model (*tpmodel*), and Fourier plane combination after deconvolution (*feather*, our adopted method in the PHANGS–ALMA pipeline). All images have the same $10''$ beam, and the color scale is fixed across images in order to allow direct comparison. The white circle in the bottom left of each panel shows the beam.

power information is crucial, though we again caution that we have created a scenario that made these effects severe.

To the eye, all the short-spacing corrected images appear fairly similar. This suggests that all three SSC methods generally do a good job of recovering the input sky model, with minor differences. However, when examined in detail, there are differences between the three results. In the rest of this appendix, we quantify the differences among the different SSC results, using three metrics.

We quantify the results for each SSC method using three measurements: the “image fidelity,” the fraction of the total model flux recovered, and the difference in peak intensity between the output and the model.

We calculate the image fidelity at each pixel defined, e.g., following ALMA memo 386 (Pety et al. 2001a), as:

$$\text{fidelity} = \frac{|\text{input model}|}{|\text{input model} - \text{output image}|}. \quad (\text{D1})$$

Thus, fidelity of 10 means that the difference between the image and the model is 9%, fidelity of 100 corresponds to a difference of $\sim 1\%$, and fidelity of 1 corresponds to a 50%

difference. More generally, the higher the fidelity, the better the output image matches the model. In practice, we compute the median fidelity over the entire image in order to quantify the overall quality of each SSC method.

The top panel of Figure 33 shows the median image fidelity as a function of the size of the CO disk for the 64 PHANGS–ALMA targets in our sample. We pick CO disk size as our independent variable because we expect targets with more extended emission, i.e., a larger CO disk size, to show more spatial filtering and thus be more dependent on the SSC. We computed the size as the diameter that contains all the nonzero pixels in the model image.

As expected, the 7 m only images show extremely low median fidelity, 1.1 with rms scatter 0.1 across the sample. This corresponds to the output image differing by $\sim 50\%$ from the input image, on average. This is consistent with the visual appearance of strong spatial filtering seen when comparing the top left and top middle panels in Figures 31 and 32. The interferometer misses the large-scale structure that contributes much of the flux in the model.

All three of the SSC methods show much higher median fidelity than the interferometer data alone, with average values

Figure 32. Results of short-spacing correction tests for NGC 4303. Similar to Figure 31, but for NGC 4303.

in the range 6.1–7.9). This implies that the reproduced images are consistent with the model within a $\sim 10\%$ – 15% accuracy. They show high scatter from target to target, however. In general, small targets show lower fidelity, even after SSC correction, compared to large targets.

The middle panel of Figure 33 shows the difference in total flux between the output image and the model as a function of the CO disk size. We define:

$$\text{total flux difference} = \frac{\text{sum of input model} - \text{sum of output image}}{\text{sum of input model}}. \quad (\text{D2})$$

In this case, 0.0 corresponds to a perfect match between input and output images.

On average, the 7 m only images miss $\sim 80\%$ of the flux. The smallest sources show better flux recovery in these images, but for virtually every source with CO disk diameter $\gtrsim 75''$, roughly 80% of the flux is missed. Recall that our test cases have discarded velocity information and so represent a worst case. In the actual PHANGS–ALMA data, the 7 m array recovers 50%–80% of the flux observed by the single dish (Figure 13) and some of the missed emission appears to be due to deconvolution effects rather than spatial filtering (Appendix C and Section 8.3).

The bottom panel in Figure 33 shows the difference in peak intensity between the output image and model. Again, we

define:

$$\text{peak intensity difference} = \frac{\text{peak in input model} - \text{peak in output image}}{\text{peak in input model}}. \quad (\text{D3})$$

This statistic captures the ability of the output image to capture the brightest emission in the image. On average, the 7 m only imaging recovers a peak flux 35% lower than in the model. Almost all 7 m images show a depressed peak flux. This shows that including total power data can be crucial even for studying compact sources. This will be especially true when, e.g., as is the case for galactic nuclei, these sources are surrounded by extended diffuse structures.

The peak fluxes of the images reconstructed by the three SSC methods agree with the model peak flux within $\sim 10\%$. The recovered peak flux does not depend on CO disk size or other galaxy parameters (e.g., total flux, peak flux, and average flux).

Combining all three metrics, we can evaluate each of the SSC methods. Before noting a few shortcomings, we emphasize that all three methods represent a marked improvement over the imaging using only the 7 m data with no SSC.

tpmodel: Images reconstructed using the **tpmodel** method show median fidelity comparable to that of those reconstructed via **feather**. They show the best overall match to the data in peak intensity. However, the middle panel of Figure 33 shows that the **tpmodel** method does tend to

Figure 33. *Measurements of the effectiveness of short-spacing corrections.* These plots show results for simulated imaging of PHANGS-ALMA data using three short-spacing correction schemes (tp2vis, tpmodel, and feather) and imaging without any short-spacing correction (“7 m only”). Each point shows a result for one reconstruction algorithm applied to simulated images based on one PHANGS-ALMA galaxy. The color of each point indicates the algorithm used for short-spacing correction. The points are sorted by the diameter of CO emission from that galaxy. The top row shows median image fidelity (Equation (D1)). The middle row shows the fraction of flux missed in the output image compared to the input model; e.g., 0.0 means no missing flux, while 0.25 means 25% of the flux is missed in the output. The bottom row shows the fractional difference in peak intensity between the output image and the model; again, 0.25 indicates a 25% lower peak intensity in the output compared to the model. The dashed lines indicate the mean performance of each method for each metric. The legend reports the mean and standard deviation for each method and metric.

recover slightly too much flux compared to the model. The total flux is overestimated by $\sim 6\%$ on average, and by as much as $\sim 25\%$ in the most extreme cases. This can be attributed to the fact that the total power data input as a model have a much larger beam size than does the CLEAN product, which creates an extended artifact surrounding strong peaks.⁵⁴ Some groups have adopted iterative strategies to address this, e.g., using feathered data from a previous iteration of clean as a model (e.g., Bolatto et al. 2013; Leroy et al. 2015). We experimented with these methods and did not find a stable, general approach, but this might represent an interesting future direction.

tp2vis: Images reconstructed using tp2vis show a high median fidelity, but with a dispersion larger than those of the other two SSC methods. The middle panel also shows that, with tp2vis, one tends to underestimate the total flux in our simulations by $\sim 22\%$ on average, and by up to $\sim 50\%$ in the most extreme cases. We find that tp2vis does a good job of recovering the peak intensity, comparable to feather and slightly worse than the tpmodel approach.

We interpret the large scatter to mean that tp2vis may require some fine-tuning of parameters for each source. Given the computing requirements for one round of imaging, as well as the fact that tp2vis did not offer a clear improvement over the other algorithms, it was not realistic to apply this fine-tuning to the current round of PHANGS-ALMA imaging. This might represent a useful future direction.

feather: Images reconstructed using feather yield a high median fidelity, ~ 7.9 , comparable to tpmodel and with a similar scatter. On average, feather recovers the total flux with accuracy similar to that of tpmodel and somewhat better than that of tp2vis. It shows lower scatter in total flux recovery than either of the other two methods. Feather tends to recover peak intensities $\sim 7\%$ too low, similar to tp2vis and slightly worse than tpmodel.

Summary: In summary, all three SSC methods represent a marked improvement over using only the 7 m data for these cases. They yield results consistent with one another at the $\sim 10\%$ level. For PHANGS-ALMA, we ultimately utilized feather because it is stable and simple with consistent performance across the sample. For ALMA, which has good

⁵⁴ <https://github.com/teuben/dc2019/blob/master/talks/Kauffmann.pdf>

overall flux calibration and a consistent flux calibration scheme for both the interferometer and total power, *feather* has the additional advantage of not requiring additional human supervision or intervention.

On average, a feathered image in our experiment has a $\sim 5\%$ bias in recovering the total flux and a $\sim 7\%$ underestimate of the peak flux. The median deviation from the input image is $\sim 12\%$, based on the image fidelity calculation. As emphasized above, these calculations significantly overstate the uncertainty, because we discard velocity information (Section 6). Still, they suggest a 5%–10% uncertainty associated with image reconstruction. This will be comparable to the uncertainty due to calibration uncertainty and thermal noise in many cases, and highlights the need for continued work on this topic.

ORCID iDs

Adam K. Leroy <https://orcid.org/0000-0002-2545-1700>
 Annie Hughes <https://orcid.org/0000-0002-9181-1161>
 Daizhong Liu <https://orcid.org/0000-0001-9773-7479>
 Jérôme Pety <https://orcid.org/0000-0003-3061-6546>
 Erik Rosolowsky <https://orcid.org/0000-0002-5204-2259>
 Toshiaki Saito <https://orcid.org/0000-0002-2501-9328>
 Eva Schinnerer <https://orcid.org/0000-0002-3933-7677>
 Antonio Usero <https://orcid.org/0000-0003-1242-505X>
 Christopher M. Faesi <https://orcid.org/0000-0001-5310-467X>
 Cinthya N. Herrera <https://orcid.org/0000-0001-6405-0785>
 Mélanie Chevance <https://orcid.org/0000-0002-5635-5180>
 Alexander P. S. Hygate <https://orcid.org/0000-0002-6488-471X>
 Amanda A. Kepley <https://orcid.org/0000-0002-3227-4917>
 Eric W. Koch <https://orcid.org/0000-0001-9605-780X>
 Miguel Querejeta <https://orcid.org/0000-0002-0472-1011>
 Christine D. Wilson <https://orcid.org/0000-0001-5817-0991>
 Gagandeep S. Anand <https://orcid.org/0000-0002-5259-2314>
 Ashley Barnes <https://orcid.org/0000-0003-0410-4504>
 Francesco Belfiore <https://orcid.org/0000-0002-2545-5752>
 Ivana Bešlić <https://orcid.org/0000-0003-0583-7363>
 Frank Bigiel <https://orcid.org/0000-0003-0166-9745>
 Guillermo A. Blanc <https://orcid.org/0000-0003-4218-3944>
 Alberto D. Bolatto <https://orcid.org/0000-0002-5480-5686>
 Mèdère Boquien <https://orcid.org/0000-0003-0946-6176>
 Yixian Cao <https://orcid.org/0000-0001-5301-1326>
 Rupali Chandar <https://orcid.org/0000-0003-0085-4623>
 Jérémy Chastenet <https://orcid.org/0000-0002-5235-5589>
 I-Da Chiang <https://orcid.org/0000-0003-2551-7148>
 Enrico Congiu <https://orcid.org/0000-0002-8549-4083>
 Daniel A. Dale <https://orcid.org/0000-0002-5782-9093>
 Sinan Deger <https://orcid.org/0000-0003-1943-723X>
 Jakob S. den Brok <https://orcid.org/0000-0002-8760-6157>
 Cosima Eibensteiner <https://orcid.org/0000-0002-1185-2810>
 Eric Emsellem <https://orcid.org/0000-0002-6155-7166>
 Axel García-Rodríguez <https://orcid.org/0000-0002-0697-0177>
 Simon C. O. Glover <https://orcid.org/0000-0001-6708-1317>
 Kathryn Grasha <https://orcid.org/0000-0002-3247-5321>
 Brent Groves <https://orcid.org/0000-0002-9768-0246>
 Jonathan D. Henshaw <https://orcid.org/0000-0001-9656-7682>
 María J. Jiménez Donaire <https://orcid.org/0000-0002-9165-8080>

Jaeyeon Kim <https://orcid.org/0000-0002-0432-6847>
 Ralf S. Klessen <https://orcid.org/0000-0002-0560-3172>
 Kathryn Kreckel <https://orcid.org/0000-0001-6551-3091>
 J. M. Diederik Kruijssen <https://orcid.org/0000-0002-8804-0212>
 Kirsten L. Larson <https://orcid.org/0000-0003-3917-6460>
 Janice C. Lee <https://orcid.org/0000-0002-2278-9407>
 Ness Mayker <https://orcid.org/0000-0002-5993-6685>
 Rebecca McElroy <https://orcid.org/0000-0003-2290-7060>
 Sharon E. Meidt <https://orcid.org/0000-0002-6118-4048>
 Angus Mok <https://orcid.org/0000-0001-7413-7534>
 Hsi-An Pan <https://orcid.org/0000-0002-1370-6964>
 Johannes Puschner <https://orcid.org/0000-0003-1111-3951>
 Alessandro Razza <https://orcid.org/0000-0001-7876-1713>
 Patricia Sánchez-BI'azquez <https://orcid.org/0000-0003-0651-0098>
 Karin M. Sandstrom <https://orcid.org/0000-0002-4378-8534>
 Amy Sardone <https://orcid.org/0000-0002-5783-145X>
 Jiayi Sun <https://orcid.org/0000-0003-0378-4667>
 David A. Thilker <https://orcid.org/0000-0002-8528-7340>
 Jordan A. Turner <https://orcid.org/0000-0003-2261-5746>
 Dyas Utomo <https://orcid.org/0000-0003-4161-2639>
 Elizabeth J. Watkins <https://orcid.org/0000-0002-7365-5791>
 Thomas G. Williams <https://orcid.org/0000-0002-0012-2142>

References

- Anand, G. S., Lee, J. C., VanDyk, S. D., et al. 2021, *MNRAS*, **501**, 3621
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Bajaja, E., & van Albada, G. D. 1979, *A&A*, **75**, 251
 Bolatto, A. D., Leroy, A. K., Rosolowsky, E., Walter, F., & Blitz, L. 2008, *ApJ*, **686**, 948
 Bolatto, A. D., Wolfire, M., & Leroy, A. K. 2013, *ARA&A*, **51**, 207
 Chevance, M., Kruijssen, J. M. D., Hygate, A. P. S., et al. 2020, *MNRAS*, **493**, 2872
 Colombo, D., Meidt, S. E., Schinnerer, E., et al. 2014, *ApJ*, **784**, 4
 Cornwell, T. J. 2008, *ISTSP*, **2**, 793
 Cotton, W. D. 2017, *PASP*, **129**, 094501
 Dirienzo, W. J., Brogan, C., Indebetouw, R., et al. 2015, *AJ*, **150**, 159
 Helfer, T. T., Vogel, S. N., Lugten, J. B., & Teuben, P. J. 2002, *PASP*, **114**, 350
 Henshaw, J. D., Kruijssen, J. M. D., Longmore, S. N., et al. 2020, *NatAs*, **4**, 1064
 Herrera, C. N., Pety, J., Hughes, A., et al. 2020, *A&A*, **634**, A121
 Heyer, M., Krawczyk, C., Duval, J., & Jackson, J. M. 2009, *ApJ*, **699**, 1092
 Heyer, M. H., Carpenter, J. M., & Snell, R. L. 2001, *ApJ*, **551**, 852
 Högbom, J. A. 1974, *A&AS*, **15**, 417
 Hughes, A., Meidt, S. E., Colombo, D., et al. 2013, *ApJ*, **779**, 46
 Kepley, A. A., Tsutsumi, T., Brogan, C. L., et al. 2020, *PASP*, **132**, 024505
 Kerr, A. R., Pan, S.-K., Claude, S. M. X., et al. 2014, *ITST*, **4**, 201
 Koch, E., Rosolowsky, E., & Leroy, A. K. 2018a, *RNAAS*, **2**, 220
 Koch, E. W., Rosolowsky, E. W., Lockman, F. J., et al. 2018b, *MNRAS*, **479**, 2505
 Koda, J., Teuben, P., Sawada, T., Plunkett, A., & Fomalont, E. 2019, *PASP*, **131**, 054505
 Landsman, W. B. 1993, in ASP Conf. Ser. 52, Astronomical Data Analysis Software and Systems II, ed. R. J. Hanisch, R. J. V. Brissenden, & J. Barnes (San Francisco, CA: ASP), 246
 Lang, P., Meidt, S. E., Rosolowsky, E., et al. 2020, *ApJ*, **897**, 122
 Lay, O. P., Carlstrom, J. E., Hills, R. E., & Phillips, T. G. 1994, *ApJL*, **434**, L75
 Leroy, A. K., Bolatto, A., Bot, C., et al. 2009, *ApJ*, **702**, 352
 Leroy, A. K., Bolatto, A. D., Ostriker, E. C., et al. 2015, *ApJ*, **801**, 25

- Leroy, A. K., Hughes, A., Schrub, A., et al. 2016, [ApJ](#), **831**, 16
- Leroy, A. K., Schinnerer, E., Hughes, A., et al. 2021, *ApJS*, submitted (arXiv:2104.07739)
- Makarov, D., Prugniel, P., Terekhova, N., Courtois, H., & Vauglin, I. 2014, [A&A](#), **570**, A13
- Mangum, J. G., Emerson, D. T., & Greisen, E. W. 2007, [A&A](#), **474**, 679
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in ASP Conf. Ser. 376, *Astronomical Data Analysis Software and Systems XVI*, ed. R. A. Shaw, F. Hill, & D. J. Bell (San Francisco, CA: ASP), 127
- McQuinn, K. B. W., Skillman, E. D., Dolphin, A. E., Berg, D., & Kennicutt, R. 2017, [AJ](#), **154**, 51
- Oliphant, T. E. 2006, *A Guide to NumPy*, Vol. 1 (USA: Trelgol Publishing)
- Pardo, J. R., Cernicharo, J., & Serabyn, E. 2001, [ITAP](#), **49**, 1683
- Paturel, G., Petit, C., Prugniel, P., et al. 2003, [A&A](#), **412**, 45
- Penzias, A. A., & Burrows, C. A. 1973, [ARA&A](#), **11**, 51
- Pety, J. 2005, in SF2A-2005: Semaine de l'Astrophysique Française, ed. F. Casoli et al. (Les Ulis: EDP Sciences), 721
- Pety, J., Gueth, F., & Guilloteau, S. 2001a, ALMA+ACA Simulation Tools, Alma Memo 386, IRAM
- Pety, J., Gueth, F., & Guilloteau, S. 2001b, Impact of ACA on the Wide-Field Imaging Capabilities of ALMA, Alma Memo 398, IRAM
- Pety, J., Schinnerer, E., Leroy, A. K., et al. 2013, [ApJ](#), **779**, 43
- R Core Team 2015, *R: A Language and Environment for Statistical Computing* (Vienna: R Foundation for Statistical Computing)
- Remijan, A. J., Markwick-Kemper, A., & ALMA Working Group on Spectral Line Frequencies 2007, AAS Meeting, **211**, 132.11
- Rosolowsky, E., Hughes, A., Leroy, A. K., et al. 2021, [MNRAS](#), **502**, 1218
- Rosolowsky, E., & Leroy, A. 2006, [PASP](#), **118**, 590
- Schinnerer, E., Hughes, A., Leroy, A., et al. 2019, [ApJ](#), **887**, 49
- Schruba, A., Leroy, A. K., Walter, F., Sandstrom, K., & Rosolowsky, E. 2010, [ApJ](#), **722**, 1699
- Shaya, E. J., Tully, R. B., Hoffman, Y., & Pomarède, D. 2017, [ApJ](#), **850**, 207
- Stanimirovic, S., Staveley-Smith, L., Dickey, J. M., Sault, R. J., & Snowden, S. L. 1999, [MNRAS](#), **302**, 417
- Sun, J., Leroy, A. K., Schinnerer, E., et al. 2020, [ApJL](#), **901**, L8
- Sun, J., Leroy, A. K., Schrub, A., et al. 2018, [ApJ](#), **860**, 172
- Teague, R., & Foreman-Mackey, D. 2018, [RNAAS](#), **2**, 173
- Tsutsumi, T., Morita, K.-I., Hasegawa, T., & Pety, J. 2004, Wide-Field Imaging of ALMA with the Atacama Compact Array: Imaging Simulations, Alma Memo 488
- Tully, R. B., Courtois, H. M., & Sorce, J. G. 2016, [AJ](#), **152**, 50
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, [NatMe](#), **17**, 261