Universiteit Leiden
The Netherlands

exoALMA. IX. Regularized maximum likelihood imaging of non-Keplerian features

Zawadzki, B.; Czekala, I.; Galloway-Sprietsma, M.; Bae, J.; Barraza-Alfaro, M.; Benisty, M.; ... ; Yoshida, T.C.

Citation
Zawadzki, B., Czekala, I., Galloway-Sprietsma, M., Bae, J., Barraza-Alfaro, M., Benisty, M., ... Yoshida, T. C. (2025). exoALMA. IX. Regularized maximum likelihood imaging of non-Keplerian features. *Astrophysical Journal Letters*, *984*(1). doi:10.3847/2041-8213/adc434

Version:	Publisher's Version
License:	Creative Commons CC BY 4.0 license
Downloaded from:	https://hdl.handle.net/1887/4290586

Note: To cite this publication please use the final published version (if applicable).

The Astrophysical Journal Letters, 984:L14 (15pp), 2025 May 1 https://doi.org/10.3847/2041-8213/adc434
© 2025. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

exoALMA. IX. Regularized Maximum Likelihood Imaging of Non-Keplerian Features

Brianna Zawadzki[1], Ian Czekala[2,3], Maria Galloway-Sprietsma[4], Jaehan Bae[4], Marcelo Barraza-Alfaro[5], Myriam Benisty[6,7,8], Gianni Cataldi[9], Pietro Curone[10,11], Stefano Facchini[10], Daniele Fasano[7,8], Mario Flock[6], Misato Fukagawa[9], Himanshi Garg[12], Cassandra Hall[13,14,15], Thomas Hilder[12], Jane Huang[16], John D. Ilee[17], Andrea Isella[18,19], Andrés F. Izquierdo[4,20,21,22], Kazuhiro Kanagawa[23], Geoffroy Lesur[7], Cristiano Longarini[10,24], Ryan A. Loomis[25], Ryuta Orihara[23], Christophe Pinte[7,12], Daniel J. Price[12], Giovanni Rosotti[10], Jochen Stadler[7,8], Richard Teague[5], Hsi-Wei Yen[26], Gaylor Wafflard-Fernandez[7], David J. Wilner[27], Andrew J. Winter[8], Lisa Wölfer[5], and Tomohiro C. Yoshida[9,28]

[1] Department of Astronomy, Van Vleck Observatory, Wesleyan University, 96 Foss Hill Drive, Middletown, CT 06459, USA; bzawadzki@wesleyan.edu
[2] School of Physics & Astronomy, University of St. Andrews, North Haugh, St. Andrews KY16 9SS, UK
[3] Centre for Exoplanet Science, University of St. Andrews, North Haugh, St. Andrews, KY16 9SS, UK
[4] Department of Astronomy, University of Florida, Gainesville, FL 32611, USA
[5] Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[6] Max-Planck Institute for Astronomy (MPIA), Königstuhl 17, D-69117 Heidelberg, Germany
[7] Université Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France
[8] Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Laboratoire Lagrange, 06304, France
[9] National Astronomical Observatory of Japan, Osawa 2-21-1, Mitaka, Tokyo 181-8588, Japan
[10] Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy
[11] Departamento de Astronomía, Universidad de Chile, Camino El Observatorio 1515, Las Condes, Santiago, Chile
[12] School of Physics and Astronomy, Monash University, Clayton VIC 3800, Australia
[13] Department of Physics and Astronomy, The University of Georgia, Athens, GA 30602, USA
[14] Center for Simulational Physics, The University of Georgia, Athens, GA 30602, USA
[15] Institute for Artificial Intelligence, The University of Georgia, Athens, GA 30602, USA
[16] Department of Astronomy, Columbia University, 538 W. 120th Str., Pupin Hall, New York, NY, USA
[17] School of Physics and Astronomy, University of Leeds, Leeds LS2 9JT, UK
[18] Department of Physics and Astronomy, Rice University, 6100 Main St, Houston, TX 77005, USA
[19] Rice Space Institute, Rice University, 6100 Main Str., Houston, TX 77005, USA
[20] Leiden Observatory, Leiden University, P.O. Box 9513, NL-2300 RA Leiden, The Netherlands
[21] European Southern Observatory, Karl-Schwarzschild-Str. 2, D-85748 Garching bei München, Germany
[22] NASA Hubble Fellowship Program Sagan Fellow
[23] College of Science, Ibaraki University, 2-1-1 Bunkyo, Mito, Ibaraki 310-8512, Japan
[24] Institute of Astronomy, University of Cambridge, Madingley Road, CB3 0HA, Cambridge, UK
[25] National Radio Astronomy Observatory, 520 Edgemont Rd., Charlottesville, VA 22903, USA
[26] Academia Sinica Institute of Astronomy & Astrophysics, 11F of Astronomy-Mathematics Building, AS/NTU, No.1, Sec. 4, Roosevelt Rd, Taipei 10617, Taiwan
[27] Center for Astrophysics—Harvard & Smithsonian, Cambridge, MA 02138, USA
[28] Department of Astronomical Science, The Graduate University for Advanced Studies, SOKENDAI, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan
Received 2024 November 18; revised 2025 February 21; accepted 2025 March 19; published 2025 April 28

Abstract

The planet-hunting Atacama Large Millimeter/submillimeter Array (ALMA) large program exoALMA observed 15 protoplanetary disks at $\sim 0\farcs15$ angular resolution and $\sim 100\,\mathrm{m\,s^{-1}}$ spectral resolution, characterizing disk structures and kinematics in enough detail to detect non-Keplerian features (NKFs) in the gas emission. As these features are often small and low-contrast, robust imaging procedures are critical for identifying and characterizing NKFs, including determining which features may be signatures of young planets. The exoALMA collaboration employed two different imaging procedures to ensure the consistent detection of NKFs: CLEAN, the standard iterative deconvolution algorithm, and regularized maximum likelihood (RML) imaging. This Letter presents the exoALMA RML images, obtained by maximizing the likelihood of the visibility data given a model image and subject to regularizer penalties. Crucially, in the context of exoALMA, RML images serve as an independent verification of marginal features seen in the fiducial CLEAN images. However, best practices for synthesizing RML images of multichanneled (i.e., velocity-resolved) data remain undefined, as prior work on RML imaging for protoplanetary disk data has primarily addressed single-image cases. We used the open-source Python package `MPoL` to explore RML image validation methods for multichanneled data and synthesize RML images from the exoALMA observations of seven protoplanetary disks with apparent NKFs in the ^{12}CO $J = 3$–2 CLEAN images. We find that RML imaging methods independently reproduce the NKFs seen in the CLEAN images of these sources, suggesting that the NKFs are robust features rather than artifacts from a specific imaging procedure.

Unified Astronomy Thesaurus concepts: Protoplanetary disks (1300); Computational methods (1965); Planet formation (1241); Deconvolution (1910); Submillimeter astronomy (1647)

1. Introduction

The Cycle 8 Atacama Large Millimeter/submillimeter Array (ALMA) large program exoALMA has observed 15 protoplanetary disks at sufficient angular and spectral resolution to

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

search for kinematic signatures of young planets by identifying localized deviations from Keplerian velocities, or non-Keplerian features (NKFs), in the disk emission (R. Teague et al. 2025). A key objective of the exoALMA program is to make reliable detections of NKFs. Interferometers like ALMA are composed of many individual antennas and deliver data in the spatial frequency domain rather than the image domain. Reconstructing images from interferometric data is an ill-posed mathematical problem that requires assumptions to be made about the unsampled spatial frequencies. Thus, recent recommendations for kinematically detecting planets strongly suggest using multiple image synthesis techniques to image the data (Disk Dynamics Collaboration et al. 2020).

The standard ALMA pipeline uses the Common Astronomy Software Applications (CASA) package for data calibration and image synthesis. The CLEAN algorithm (especially the CASA implementation, `tclean`), is the most common method for imaging radio interferometric data and, in the absence of noise, is equivalent to a least squares fit to the visibilities (J. A. Högbom 1974; U. J. Schwarz 1978; J. P. McMullin et al. 2007; U. Rau & T. J. Cornwell 2011; CASA Team et al. 2022). CLEAN is a procedural algorithm that builds a model image by identifying the brightest pixels in the dirty image and placing a CLEAN component at the same location in the model image. Then, the CLEAN component is convolved with the dirty beam (the instrument point-spread function) and subtracted from the dirty image. At the end of the CLEANing process (e.g., after a specified number of iterations or noise threshold is reached), there are two products: the residual image and the model image. The residual image is what remains after all the dirty-beam-convolved CLEAN components have been subtracted out of the original dirty image. The model image, composed of the many CLEAN components, is typically added to the residual image and convolved with a restoring beam (the "CLEAN beam") to obtain a final cleaned image. The fiducial exoALMA images were made with CLEAN, as described in R. Teague et al. (2025), R. Loomis et al. (2025).

However, despite the fact that the ALMA pipeline and users rely primarily on CLEAN for imaging, it is ultimately a heuristic, user-dependent algorithm that attempts to build well-resolved and often complex structures from a set of point source or Gaussian components. While this typically works well enough, the hierarchical ringed substructures commonly seen in disk observations (e.g., S. M. Andrews et al. 2018) are particularly poorly matched to the basis sets of points and Gaussians. Thus, the community should consider adopting imaging methods that do not assume resolved sources like disks are represented by a collection of discrete points. One alternative to CLEAN is regularized maximum likelihood (RML) imaging, a forward-modeling approach to imaging which aims to identify the most likely set of image pixels given the data and chosen regularizers (e.g., T. J. Cornwell & K. F. Evans 1985; R. Narayan & R. Nityananda 1986; M. Cárcamo et al. 2018; Event Horizon Telescope Collaboration et al. 2019). RML imaging is an optimization-based imaging framework; the image is obtained by maximizing a likelihood function subject to regularizer penalties. The precise form of the likelihood function can vary depending on the data characteristics (e.g., noise qualities) and which regularizers are implemented. In this work, RML models are directly parameterized by the image pixels and are not composed of CLEAN-like components nor convolved with a restoring beam (which can artificially degrade the resolution of the image).

Though CLEAN and RML are distinct image synthesis methods, both have been used successfully and complimentarily for ALMA protoplanetary disk observations (e.g., M. Cárcamo et al. 2018; S. Pérez et al. 2019; M. Yamaguchi et al. 2020, 2021; S. Casassus et al. 2021; B. Zawadzki et al. 2023). In particular, M. Cárcamo et al. (2018, 2019) introduced the `GPUVMEM` package, which enabled GPU-accelerated RML imaging with a focus on ALMA protoplanetary disk imaging, establishing the use of efficient RML frameworks in the disk community. Utilizing multiple imaging methods is particularly worthwhile for testing the robustness of small or otherwise tenuous features within a disk, such as emission features likely caused by disk instabilities or embedded planets. Detecting these features in multiple independently synthesized images provides an additional layer of confidence that these features are real, as these features can be subtle and difficult to distinguish from spurious features arising from the imaging procedure rather than the data itself (as done with the Event Horizon Telescope observations of M87; e.g., A. A. Chael et al. 2018; Event Horizon Telescope Collaboration et al. 2019). For example, the imperfect subtraction of the instrument PSF during the CLEANing process could cause faint speckles to appear in the image, causing misinterpretations in a study attempting to directly image planets. Suppose an NKF appears in the source emission regardless of the imaging method used. In that case, the detected features are more likely real than an artifact stemming from a specific imaging choice. If the detected features do not appear in independently synthesized images, then extra care should be taken to validate the feature before drawing conclusions about the presence of protoplanets or other disk processes. In order to fully characterize the physical and dynamical structures of the disk at high confidence, a comprehensive approach to imaging is essential.

Artifacts may be introduced either through imperfect calibration or imaging procedures, the former of which would affect both the CLEAN and RML images. R. Loomis et al. (2025) describe steps taken to validate the exoALMA calibration process, and here we focus on improving confidence in the imaging results. We present the results of RML imaging applied to seven exoALMA sources that displayed prominent NKFs in their ^{12}CO $J = 3$–2 CLEAN images. The source selection was motivated by a preliminary visual search for kink-like substructures caused by planets, resulting in an initial sample of five disks: AA Tau, J1615, J1842, LkCa 15, and SY Cha. These are the same five sources for which C. Pinte et al. (2025) find evidence of embedded protoplanets and place constraints on the planet masses. To explore the effects of RML imaging on a wider variety of substructures, we add HD 135344B and J1604 to our sample, which both display large-scale non-Keplerian arcs (see Figure 1 in C. Pinte et al. 2025). For each of these seven sources, we present full ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 image cubes, as well as comparisons to the CLEAN images for select channels of interest.

2. Source Selection and Data Processing

The exoALMA large program obtained data of 15 large and bright protoplanetary disks at moderate angular (maximum baselines up to 3697 m) and high spectral (up to 30.5 kHz, 26 m s^{-1}) resolution (R. Loomis et al. 2025;

R. Teague et al. 2025). Here, we use CLEAN image cubes with 100 m s^{-1} channel spacings and 0.″15 restoring beams as the fiducial image products for comparison with the RML image products; this is equivalent to the fiducial set of exoALMA images for ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2, but with a finer channel spacing than the CS $J = 7$–6 fiducial set of images (which has 200 m s^{-1} spacings). The full data calibration and CLEAN imaging procedures are detailed in R. Loomis et al. (2025).

In this Letter, we focus on select exoALMA targets that show clear NKFs in the ^{12}CO $J = 3$–2 emission imaged with `tclean`: protoplanetary disks around AA Tau, HD 135344B, J1604, J1615, J1842, LkCa 15, and SY Cha. For an in-depth analysis of the NKFs in the exoALMA sample, see C. Pinte et al. (2025). To prepare the data for RML imaging, we first transformed each calibrated measurement set (i.e., the set of (u,v) points, complex visibilities, and visibility weights at each frequency) using the `CASA cvel2` task such that the data were in an LSRK reference frame. This must be done explicitly prior to exporting the data for RML imaging, as the reference frame correction is normally built into the `tclean` procedure. The data were regridded to a channel width of 0.1 km s^{-1}; however, the regridding process is not necessary for RML imaging. We regridded the data to enable a direct comparison to the fiducial CLEAN images.

We use the RML imaging package `MPoL` for all imaging (I. Czekala et al. 2023, 2025; B. Zawadzki et al. 2023). We exported each of the regridded measurement sets (both the continuum-subtracted and non-continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 measurement sets for each disk) as arrays of complex visibilities for use with `MPoL`, and averaged the ungridded visibility data to grid cells in the visibility domain before imaging (equivalent to uniform weighting). The grid cells are specified in the image domain; in this study, all RML images are 1024×1024 pixels, with each pixel measuring 0.″0125 across (half the size of the fiducial CLEAN pixels). This pixel scaling was chosen to ensure that the data with the highest angular resolutions (up to 0.″06) still have several pixels across a resolution element while still capturing the full extent of the disk emission. We then defined a corresponding 1024×1024 Fourier grid and averaged the ungridded visibility data to the Fourier cells using a simple weighted average. We have verified that this gridding procedure does not introduce interpolation artifacts to the images at the signal-to-noise ratio (SNR) of the data.

3. RML Techniques for Image Cubes

RML imaging aims to find the set of visibilities that maximizes the likelihood function

$$p(\boldsymbol{D}|\boldsymbol{I}), \tag{1}$$

where $\boldsymbol{D}$ is a set of visibility data and $\boldsymbol{I}$ is a model image. In practice, we compute the negative log likelihood for computational efficiency. Assuming that the noise is normally distributed with standard deviation σ and is uncorrelated across baselines, the negative natural logarithm of the likelihood function is

$$-\ln p(\boldsymbol{D}|\boldsymbol{I}) = N_{\mathrm{D}} \ln(\sqrt{2\pi}\sigma) + \frac{1}{2}\sum_{i}^{N_{\mathrm{D}}} \left| \frac{D_i - V_i(\boldsymbol{I})}{\sigma_i} \right|^2, \tag{2}$$

where N_D is the number of complex visibilities in the data set, D_i is a measured complex visibility at a (u_i, v_i) point, and V_i is the predicted value of the model visibilities for the same (u_i, v_i) generated from the model image. The rightmost term is equivalent to $1/2$ of the χ^2 statistic, and the negative log likelihood can be written as

$$L_{\mathrm{nll}}(\boldsymbol{I}) = -\ln p(\boldsymbol{D}|\boldsymbol{I}) = \frac{1}{2}\chi^2(\boldsymbol{D}|\boldsymbol{I}). \tag{3}$$

The $L_{\mathrm{nll}}(\boldsymbol{I})$ notation is adapted from the machine learning community, where it is common to identify a function that can be minimized in order to obtain optimal model parameter values, known as a loss function (C. M. Bishop 2006). Though minimizing $L_{\mathrm{nll}}(\boldsymbol{I})$ is equivalent to maximizing the likelihood function, the resulting image may not be optimal due to the incomplete (u, v) sampling of the data. In other words, the ill-posed and underdetermined nature of radio interferometric imaging means that the maximum likelihood image cannot be uniquely determined. An image product can be improved from this base maximum likelihood image by incorporating one or more regularizers into the loss function; we adopt the same basic loss function as in B. Zawadzki et al. (2023, hereafter Z23), using a combination of entropy, sparsity, and total squared variation (TSV; K. Kuramochi et al. 2018) regularization.

3.1. The Loss Function

The maximum entropy loss is defined as

$$L_{\mathrm{ent}} = \frac{1}{\zeta}\sum_i I_i \ln I_i, \tag{4}$$

where I_i is the intensity of an image pixel. Here, ζ is a normalization factor, which we set equal to the total flux of the CLEAN image. Maximum entropy regularization restricts pixel values to positive nonzero values and favors uniformity in the image, making it useful for identifying emission features (S. F. Gull & G. J. Daniell 1978; J. A. Högbom 1979; R. Narayan & R. Nityananda 1986).

The sparsity loss is defined as

$$L_{\mathrm{spa}} = \sum_i |I_i| \tag{5}$$

and is derived from the least absolute shrinkage and selection operator (lasso; R. Tibshirani 1996). Sparsity regularization uses the L_1 norm to suppress the amplitudes of low-intensity pixels, promoting an image that is a sparse collection of nonzero valued pixels. This often has the effect of greatly reducing background noise in the image and improving image resolution (M. Honma et al. 2014).

Finally, the TSV loss is defined as

$$L_{\mathrm{TSV}} = \sum_{l,m} (I_{l+1,m} - I_{l,m})^2 + (I_{l,m+1} - I_{l,m})^2, \tag{6}$$

where l and m are indices over pixels corresponding to R.A. and decl., respectively. TSV regularization takes the contiguity of image pixels into account, favoring sharp changes in intensity when needed and relatively smooth areas elsewhere.

Each of the above regularizers has an adjustable coefficient λ, which determines the strength of that regularizer. Since each λ affects the values of the model parameters ($\boldsymbol{I}$) without itself

being a free parameter in the model, it is known as a hyperparameter. The total loss function is the sum of the likelihood function and all regularizing terms,

$$L(\boldsymbol{I}) = L_{\rm nll}(\boldsymbol{I}) + \lambda_{\rm ent} L_{\rm ent}(\boldsymbol{I}) + \lambda_{\rm spa} L_{\rm spa}(\boldsymbol{I}) + \lambda_{\rm TSV} L_{\rm TSV}(\boldsymbol{I}). \quad (7)$$

3.2. Cross Validation for Multichanneled Data

Cross-validation methods are well-established in astronomy as a quantitative procedure for obtaining hyperparameter values (e.g., K. Akiyama et al. 2017a, 2017b; M. Aizawa et al. 2020; M. Yamaguchi et al. 2020, 2021, 2024; B. Zawadzki et al. 2023). Following the suggestions in Z23, we use random cell cross validation (CV) to identify the optimal hyperparameters ($\lambda_{\rm ent}$, $\lambda_{\rm spa}$, and $\lambda_{\rm TSV}$) for imaging. Random cell CV works by partitioning visibility grid cells into K pairs of training and testing sets randomly, drawing cells without replacement with the minor exception that grid cells with the highest 1% of gridded weight values are included in each subset in order to avoid numerical instabilities that delay or prevent convergence. After defining the loss function, an RML model is generated using the training data. This model is then compared to the withheld testing data to quantify the predictive power of the RML model and obtain a CV score,

$$\mathrm{CV} = \sum_{k}^{K} \chi^2(\boldsymbol{D}_k | \boldsymbol{I}_{\rm train}), \quad (8)$$

where $\boldsymbol{D}_k$ is the test data set, and $\boldsymbol{I}_{\rm train}$ is the model image that minimizes the loss function for the training data $\boldsymbol{D} - \boldsymbol{D}_k$. We use $K = 10$, which has been shown to effectively balance bias and variance in the parameter error estimates (L. Breiman & P. Spector 1992; R. Kohavi 1995; A. M. Molinaro et al. 2005).

Previous work on RML imaging for protoplanetary disk observations focused primarily on continuum images (e.g., M. Cárcamo et al. 2018; S. Casassus & S. Pérez 2019; S. Pérez et al. 2019, 2020; M. Yamaguchi et al. 2020, 2021, 2024; B. Zawadzki et al. 2023); however, extending the application of RML methods to velocity-resolved gas observations presents a new set of challenges. While some studies have used RML methods for imaging the gas emission in protoplanetary disks (e.g., M. Cárcamo et al. 2018), only a few present velocity-resolved RML channel maps (S. Casassus et al. 2021, 2022). As a result, there is not yet a firm consensus on best practices for applying RML techniques to multichanneled data like the observations presented here. In particular, it is unclear whether CV must be performed on a channel-by-channel basis or whether CV results from a representative channel can be applied to the rest of (or at least a subsection of) the full data cube. It is also unclear whether optimal hyperparameter values are comparable for different observations or data products of the same source (e.g., ^{12}CO $J = 3$–2 versus ^{13}CO $J = 3$–2, or continuum-subtracted versus non-continuum-subtracted data). The CV process is by far the most computationally expensive component of RML imaging, so eliminating unnecessary hyperparameter tuning results in significant time and resource savings when working with a large sample like the exoALMA disks.

To test this, we performed CV on the LkCa 15 data, evaluating a grid of entropy, TSV, and sparsity hyperparameter values spanning several orders of magnitude. We included the ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 observations and tested both the continuum-subtracted and non-continuum-subtracted observations across multiple different channels sampling different parts of the velocity space. Our findings are summarized in Table 1, showing results from a subset of our 30 CV tests. In total, we tested five nonadjacent channels for each line (both continuum-subtracted and non-continuum-subtracted). These channels are presented in Figure 1 and were selected to probe different spatial extents and general morphologies of the emission.

Table 1
Top: Tested Range of Regularizer Hyperparameter Values ($\lambda_{\rm ent}$, $\lambda_{\rm spa}$, and $\lambda_{\rm TSV}$)

Tested Hyperparameter Values					
$\lambda_{\rm ent}$	[0, 8e-8, 2e-7, 8e-7, 2e-6, 8e-6, 2e-5]				
$\lambda_{\rm TSV}$	[5e-5, 1e-4, 5e-4, 1e-3, 5e-3, 1e-2]				
$\lambda_{\rm spa}$	[5e-6, 1e-5, 5e-5]				
line	contsub?	channel	$\lambda_{\rm ent}$	$\lambda_{\rm TSV}$	$\lambda_{\rm spa}$
^{12}CO $J = 3$–2	Yes	6.9 km s^{-1}	0	5e-4	1e-5
^{12}CO $J = 3$–2	Yes	6.3 km s^{-1}	0	5e-4	1e-5
^{12}CO $J = 3$–2	Yes	4.8 km s^{-1}	0	1e-4	1e-5
^{12}CO $J = 3$–2	No	6.9 km s^{-1}	0	5e-4	1e-5
^{12}CO $J = 3$–2	No	6.3 km s^{-1}	0	5e-4	1e-5
^{12}CO $J = 3$–2	No	4.8 km s^{-1}	8e-8	1e-4	1e-5
^{13}CO $J = 3$–2	Yes	6.9 km s^{-1}	0	5e-4	1e-5
^{13}CO $J = 3$–2	Yes	6.3 km s^{-1}	0	5e-4	1e-5
^{13}CO $J = 3$–2	Yes	4.8 km s^{-1}	0	5e-4	1e-5
CS $J = 7$–6	Yes	6.9 km s^{-1}	0	5e-4	1e-5
CS $J = 7$–6	Yes	6.3 km s^{-1}	0	5e-4	1e-5
CS $J = 7$–6	Yes	4.8 km s^{-1}	0	1e-3	1e-5

Note. Botton: optimal entropy, TSV, and sparsity λ values resulting from random cell CV on various LkCa 15 data. The full CV process was conducted on (1) ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data, (2) both the continuum-subtracted and non-continuum-subtracted observations, and (3) five individual channels (100 m s^{-1} channel width) across the data cubes displaying different emission morphologies, including the channel with the apparent NKF at 6.9 km s^{-1}. The optimal λ values resulting from CV show little variation across channels or lines, suggesting that CV on a single representative channel is sufficient for setting hyperparameter values for an entire data cube or multiple similar observations of the same source, greatly reducing the computational burden of RML imaging across large data sets.

The optimal hyperparameter values tended to remain nearly constant throughout a given image cube, with only slight variations from channel to channel or no variation at all. We verified that these minor variations did not significantly change the emission morphologies in the resulting images. Optimal hyperparameters resulting from the CV procedures also remained largely consistent between the continuum-subtracted and non-continuum-subtracted data. Lastly, hyperparameter values remained similar across the ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data despite often having substantial differences in emission morphology and intensity.

These results suggest that hyperparameter values are more dependent on the observational parameters (e.g., interferometer baselines and on-source integration time, corresponding to the resolution and sensitivity of the data) rather than the intensities or specific morphological features across different molecular lines. Thus, it is possible to exploit similarities between data to significantly reduce the computational burden of CV for RML hyperparameter tuning. A typical CV procedure for one combination of hyperparameters can take 30 minutes on a NVIDIA P100 GPU; with a comprehensive grid of

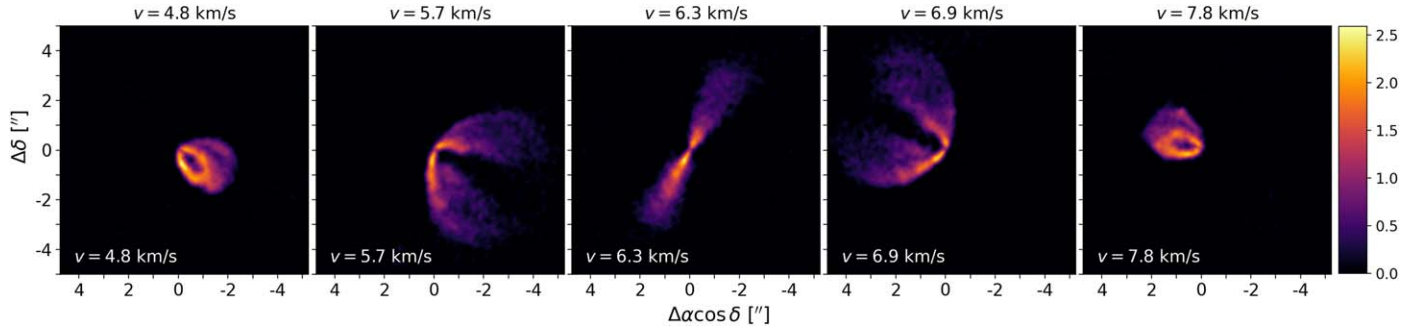

Figure 1. Five channels of LkCa 15 selected for thorough CV testing (RML images pictured). Channels were selected arbitrarily across the range of velocity space where gas emission is prominent in all three molecular lines, testing channels that varied in the spatial extent and general morphology of the emission. We intentionally include a channel with a prominent NKF at $v = 6.9$ km s^{-1}. The color bar has units of Jy arcsec^{-2}.

hyperparameter values, the process requires tens of hours (1–2 days) on a single GPU. If the CV-tuned hyperparameters showed substantial variation within and across the cubes, obtaining hyperparameter values with CV would quickly become impractical for large data sets like the exoALMA data.

Given these results, we followed the RML workflow laid out in Z23, except adapted for working with multichanneled data:

1. For each source, we selected a single channel to use for CV by examining the fiducial continuum-subtracted ^{12}CO $J = 3$–2 data CLEAN images. For sources with a spatially small NKF localized to a few channels, we used one of the channels where the NKF is visible in the fiducial CLEAN images. For sources with large-scale NKFs spanning many channels, we selected a representative channel where the NKFs are clearly seen in the fiducial CLEAN images. Following the extensive CV testing summarized in Table 1, we do not expect the choice of channel at this step to have a significant impact on the final RML images.
2. We performed tenfold CV on the continuum-subtracted ^{12}CO $J = 3$–2 data at this channel for the range of hyperparameter values ($\lambda_{\rm ent}$, $\lambda_{\rm TSV}$, and $\lambda_{\rm spa}$) presented in Table 1.
3. We applied the hyperparameter values that minimized the CV score at this channel to our loss function (Equation (8)). We used this loss function for all channels in the cube.
4. With the loss function for each source now specified with the optimized hyperparameters, we synthesized full RML image cubes for the continuum-subtracted and non-continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data.

The CV results are summarized in Table 2. Five of the seven sources yielded optimal $\lambda_{\rm ent}$ values of zero, effectively "turning off" entropy regularization and indicating that TSV and sparsity regularization alone yielded the model with the best predictive power between the training and testing data during CV. Other recent studies have also obtained high-quality protoplanetary disk images using only a combination of TSV and sparsity regularization (e.g., M. Yamaguchi et al. 2024). The two remaining disks, J1615 and J1842, favor a modest degree of entropy regularization, and each has an optimal $\lambda_{\rm ent}$ of 8×10^{-6}. For all seven disks, the optimal $\lambda_{\rm TSV}$ values range from 5×10^{-5} to 5×10^{-4}, and the optimal $\lambda_{\rm spa}$ values range from 1×10^{-5} to 5×10^{-5}.

We note that it is also possible to continue tuning hyperparameters by hand after the CV process. However, deviating from the set of hyperparameters that minimize the CV score introduces human bias, as hand-tuning is simply adjusting the hyperparameters manually to obtain a more desirable image appearance (evaluated by eye rather than a quantitative metric like a CV score; see the discussion on hyperparameter tuning in Z23). To treat the exoALMA sample more uniformly, we do not take this optional tuning step and use strictly the hyperparameter values that minimize the CV score for each source.

Table 2
Hyperparameter Values Resulting from Tenfold CV on the Continuum-subtracted ^{12}CO $J = 3$–2 Data for Each Source

Source	Channel (km s^{-1})	$\lambda_{\rm ent}$	$\lambda_{\rm TSV}$	$\lambda_{\rm spa}$
AA Tau	$v = 7.9$	0	5e-5	1e-5
HD 135344B	$v = 6.6$	0	5e-5	5e-5
J1604	$v = 4.5$	0	1e-4	5e-5
J1615	$v = 3.7$	8e-6	5e-4	5e-5
J1842	$v = 5.1$	8e-6	5e-5	5e-5
LkCa 15	$v = 6.9$	0	5e-4	1e-5
SY Cha	$v = 3.6$	0	5e-4	1e-5

Note. These values were obtained from running CV on a single channel of interest (100 m s^{-1} channel width), then applied to the rest of the cube.

We also find that the optimal hyperparameter values remain similar across different exoALMA sources; however, there are small differences between most sources (Table 2). We do not automatically apply any results from one source to another, conducting CV independently for each disk. As optimal hyperparameter values appear to be driven primarily by the (u, v) coverage and sensitivity of the observations, CV results from one target can be used as a starting point for the rest of the sample in large, relatively uniform programs like exoALMA. Given that only two disks (LkCa 15 and SY Cha) had identical optimal hyperparameters, however, we recommend performing CV at least once per source.

This apparent dependence on (u, v) coverage and sensitivity also means that the simple CV procedure used here may not be as effective for observations with less uniform (u, v) coverage or lower SNRs. For applications to such data, it may be necessary to repeat the CV process multiple times with different visibility partitioning to further validate and avoid artificially skewing the hyperparameter values based on the random partitioning into training and testing sets. In this case, it would also be useful to test the performance of CV on simulated data with comparably sparse or nonuniform (u, v)

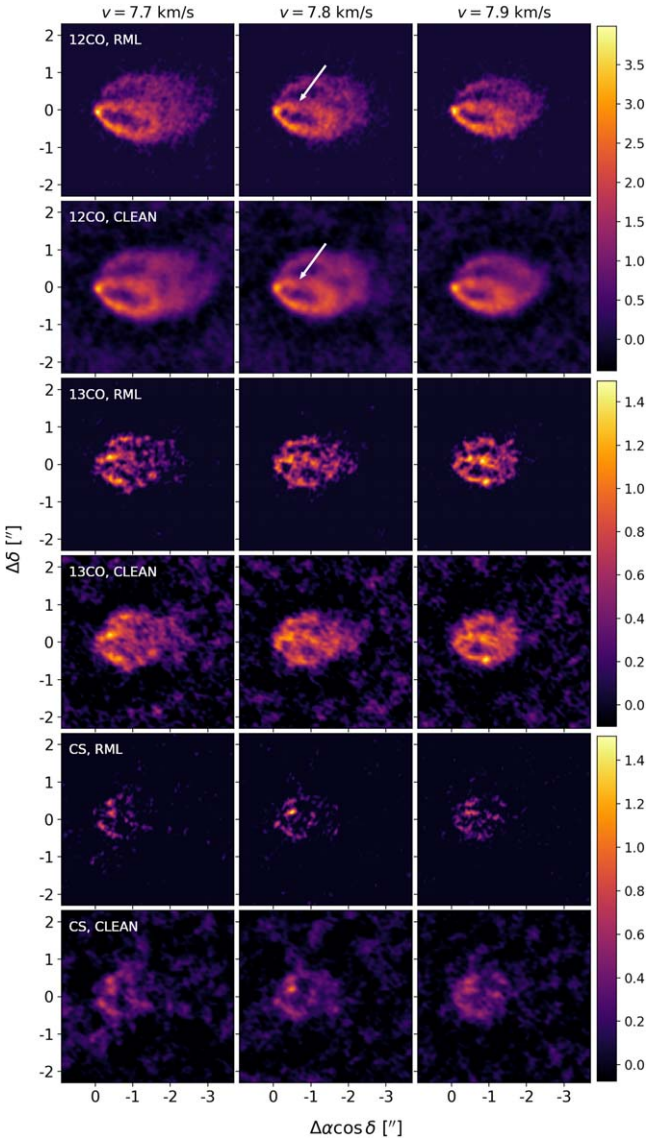

Figure 2. RML and CLEAN comparison for continuum-subtracted observations of AA Tau in ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6. Shown are three adjacent channels, centered at $v = 7.8$ km s^{-1} where the ^{12}CO $J = 3$–2 non-Keplerian feature appears most prominently. All six panels for each molecular line (the three RML panels and three CLEAN panels) are plotted on the same color scale, and the color bars for each molecule are in units of Jy arcsec^{-2}. The white arrows indicate the NKF, which can be seen in the ^{12}CO $J = 3$–2 emission; while the feature is present in all three channels shown, we only place the arrows in one RML image and one CLEAN image to highlight the location of the feature.

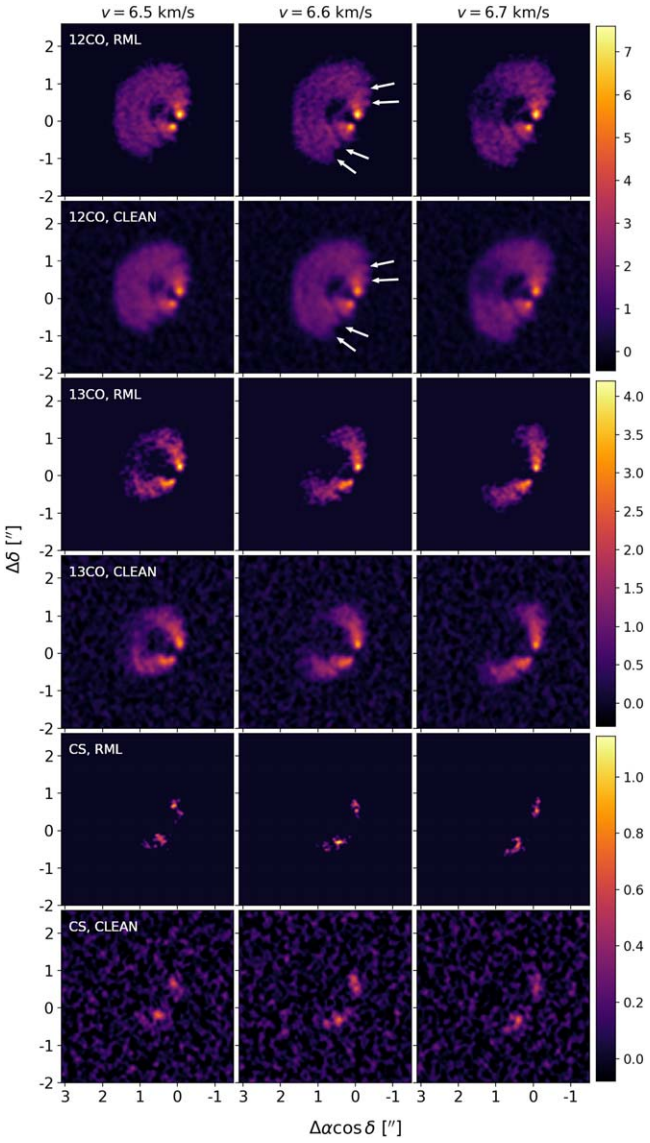

Figure 3. Same as Figure 2, but for HD 135344B. As HD 135344B displays large NKFs that persist throughout the cube, we show three adjacent and representative channels centered at $v = 6.6$ km s^{-1}.

coverage to better characterize the behavior of CV in this low sensitivity regime.

4. Results

Following the recommendations from Z23 and incorporating the new findings for spectral line data described in Section 3, we generated continuum-subtracted and non-continuum-subtracted RML image cubes of the ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 emission with 100 m s^{-1} channel spacing for seven exoALMA targets. Each source is summarized with a figure comparing the RML and CLEAN images (Figures 2–8). Each figure displays three channels, centered on a channel where the NKF(s) appear prominently in the ^{12}CO $J = 3$–2 data (see C. Pinte et al. 2025). Figure 9 displays the most prominent NKFs in the continuum-subtracted ^{12}CO $J = 3$–2 emission for each source, showing the difference between the RML and CLEAN images with a zoomed-in look at the NKFs. The RML images presented in these figures were made with the optimal hyperparameter values presented in Table 2, following the procedure described in Section 3.2.

4.1. AA Tau

The CLEAN images of AA Tau show an NKF in the ^{12}CO $J = 3$–2 emission centered at roughly $v = 7.8$ km s^{-1} and visible in multiple adjacent channels. Figure 2 shows a comparison of the RML and CLEAN images of AA Tau for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. All of the RML images in Figure 2 were made with $\lambda_{\rm ent} = 0$, $\lambda_{\rm TSV} = 5 \times 10^{-5}$, and $\lambda_{\rm spa} = 1 \times 10^{-5}$.

The RML images of AA Tau appear sharper than the corresponding CLEAN images. All the features appearing in the CLEAN images are recovered in the corresponding RML

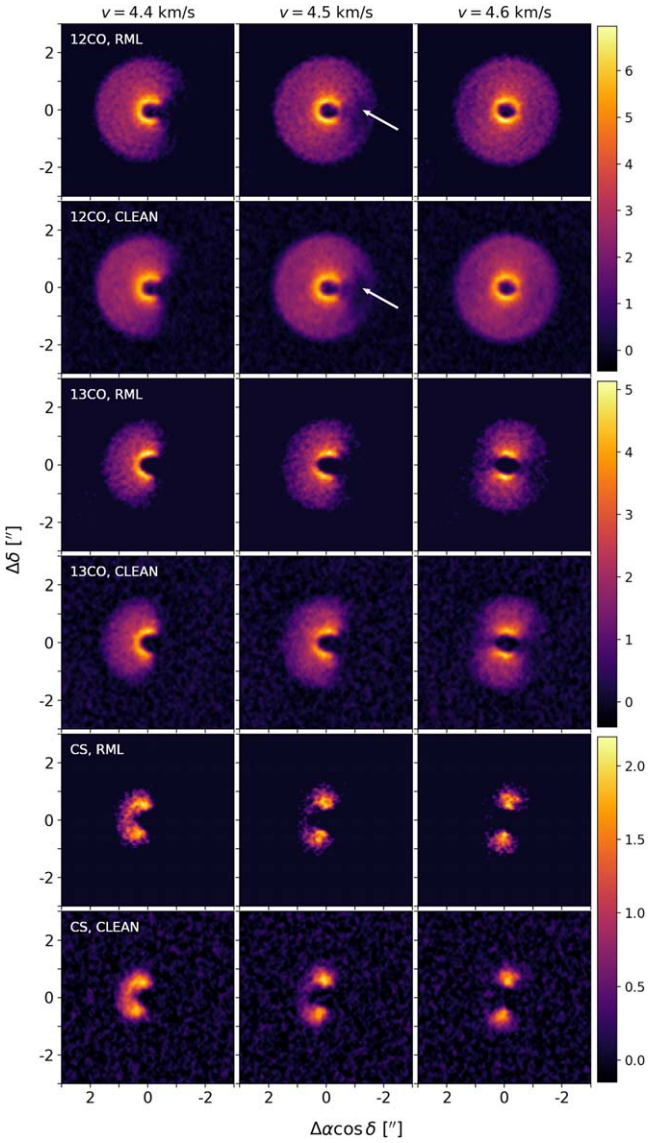

Figure 4. Same as Figure 2, but for J1604. As J1604 displays large NKFs that persist throughout the cube, we show three adjacent and representative channels centered at $v = 4.5$ km s^{-1}.

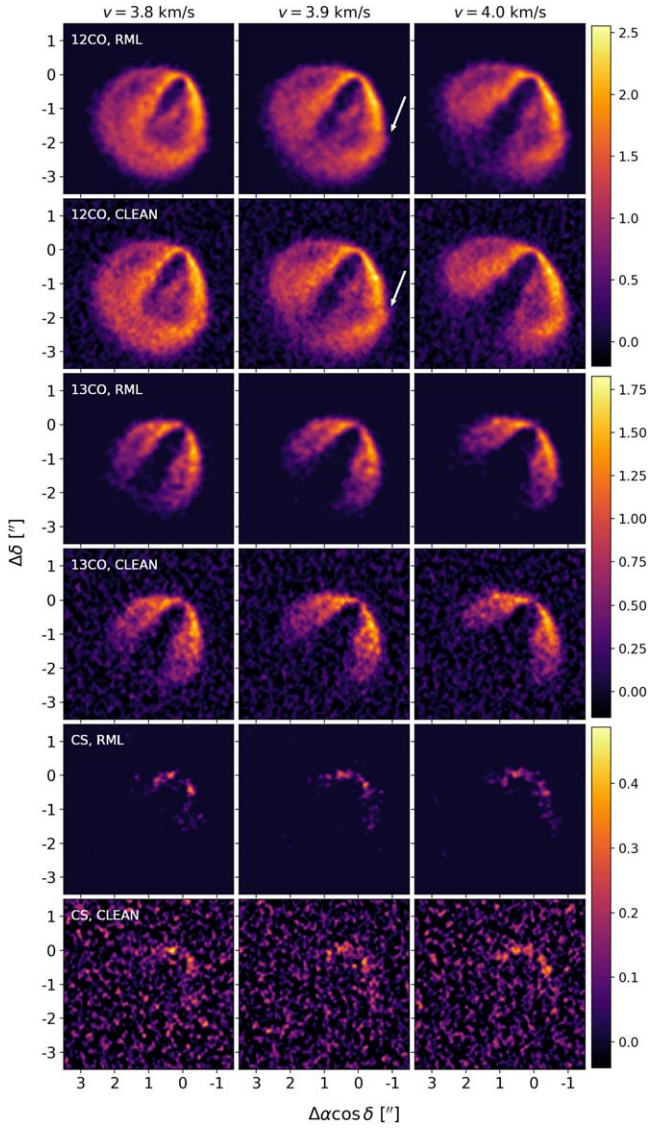

Figure 5. Same as Figure 2, but for J1615. Shown are three adjacent channels, centered at $v = 3.9$ km s^{-1} where the ^{12}CO $J = 3$–2 non-Keplerian feature appears most prominently.

images, except the RML images have more uniform background emission due to the sparsity regularization. This can also be thought of as background noise suppression through regularization. The NKF is visible in both the RML and CLEAN ^{12}CO $J = 3$–2 images and is consistent with a 2 $M_{\rm Jup}$ planet embedded in the disk (C. Pinte et al. 2025). This feature is not as apparent in the ^{13}CO $J = 3$–2 or CS $J = 7$–6 images, but there are bright clumps present in these channels at the rough location of the NKF.

4.2. HD 135344B

The CLEAN images of HD 135344B show NKFs consistently throughout the ^{12}CO $J = 3$–2 emission. Figure 3 shows a comparison of the RML and CLEAN images of HD 135344B for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. Though the NKFs in this disk are not localized to only a handful of channels like planet-driven NKFs, we select three adjacent channels centered at $v = 6.6$ km s^{-1} for illustrative purposes. All of the RML images in Figure 3 were made with $\lambda_{\rm ent} = 0$, $\lambda_{\rm TSV} = 5 \times 10^{-5}$, and $\lambda_{\rm spa} = 5 \times 10^{-5}$.

The RML images of HD 135344B appear sharper than the corresponding CLEAN images and have more uniform background emission due to the sparsity regularization. The RML images consistently reproduce the large-scale NKFs seen in the CLEAN images.

4.3. J1604

The CLEAN images of J1604 show NKFs consistently throughout the ^{12}CO $J = 3$–2 emission. Figure 4 shows a comparison of the RML and CLEAN images of J1604 for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. Though the NKFs in this disk are not localized to only a handful of channels like planet-driven NKFs, we select three adjacent channels centered at $v = 4.5$ km s^{-1} for illustrative purposes. All of the RML images in Figure 4 were made with $\lambda_{\rm ent} = 0$, $\lambda_{\rm TSV} = 1 \times 10^{-4}$, and $\lambda_{\rm spa} = 5 \times 10^{-5}$.

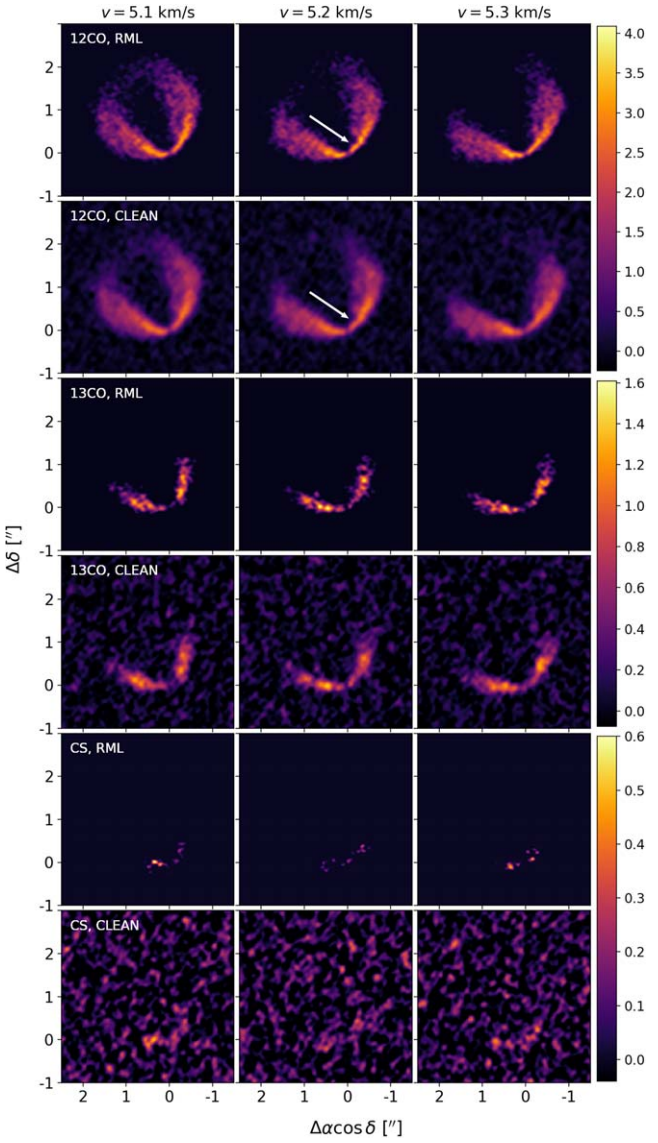

Figure 6. Same as Figure 2, but for J1842. Shown are three adjacent channels, centered at $v = 5.2$ km s^{-1} where the ^{12}CO $J = 3$–2 non-Keplerian feature appears most prominently.

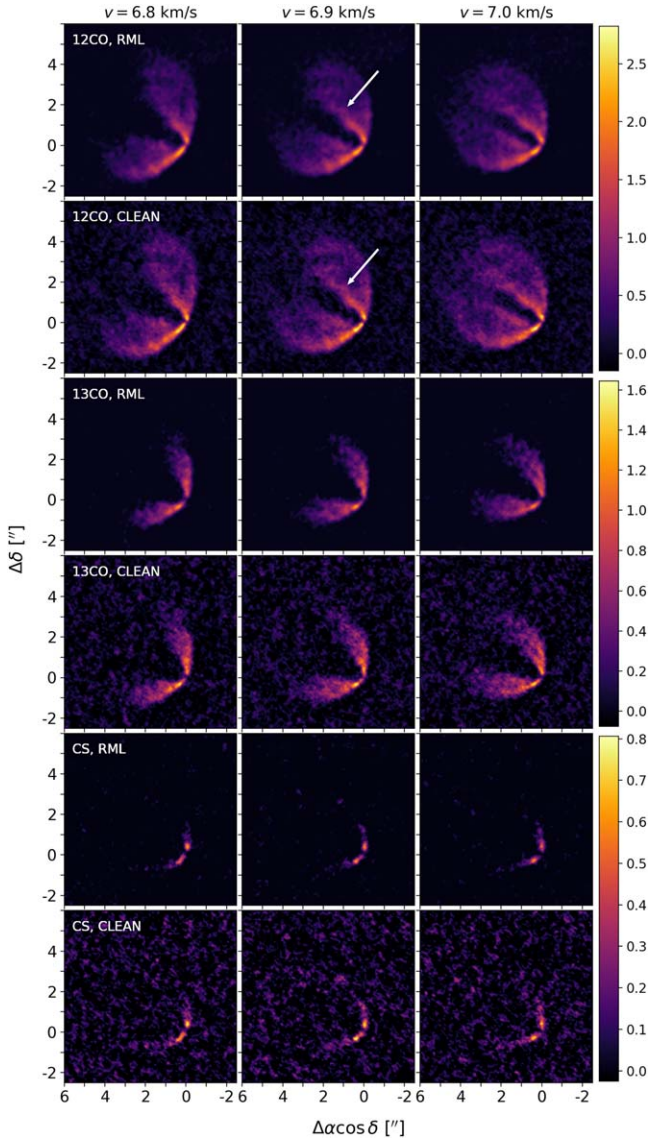

Figure 7. Same as Figure 2, but for LkCa 15. Shown are three adjacent channels, centered at $v = 6.9$ km s^{-1} where the ^{12}CO $J = 3$–2 non-Keplerian feature appears most prominently.

Here, the RML images look extremely similar to the fiducial CLEAN images, with only a small degree of added sharpness, which appears more obvious in the fainter ^{13}CO $J = 3$–2 and CS $J = 7$–6 emission. Features like the smoothness of the emission, edge sharpness, and presence of NKFs appear to be fully reproduced in the RML images. Background noise is suppressed in the RML images due to sparsity regularization.

4.4. J1615

The CLEAN images of J1615 show an NKF in the ^{12}CO $J = 3$–2 emission centered at roughly $v = 3.9$ km s^{-1}. Figure 5 shows a comparison of the RML and CLEAN images of J1615 for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. All of the RML images in Figure 5 were made with $\lambda_{\rm ent} = 8 \times 10^{-6}$, $\lambda_{\rm TSV} = 5 \times 10^{-4}$, and $\lambda_{\rm spa} = 5 \times 10^{-5}$.

While the RML images appear slightly more blurred than the CLEAN images for this target, the NKF remains visible in multiple channels. This feature is consistent with a $2\,M_{\rm Jup}$ planet embedded in the disk (C. Pinte et al. 2025). Incorporating the additional step of hand-tuning hyperparameter values after CV may result in a sharper RML image. The largest difference between the RML and CLEAN images is in the CS $J = 7$–6 emission; the disk emission stands out more prominently in the RML images than in the CLEAN images, which are comparatively noise dominated. This is due to the sparsity regularization incorporated into the RML model, which reduces the number of image pixels with nonzero values and effectively suppresses background noise in the image.

4.5. J1842

The CLEAN images of J1842 show an NKF in the ^{12}CO $J = 3$–2 emission centered at roughly $v = 5.2$ km s^{-1}. Figure 6 shows a comparison of the RML and CLEAN images of J1842 for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. All of the RML images in Figure 6 were made with $\lambda_{\rm ent} = 8 \times 10^{-6}$, $\lambda_{\rm TSV} = 5 \times 10^{-5}$, and $\lambda_{\rm spa} = 5 \times 10^{-5}$.

The RML images are sharper than the corresponding CLEAN images, and the NKF seen in the ^{12}CO $J = 3$–2 CLEAN images

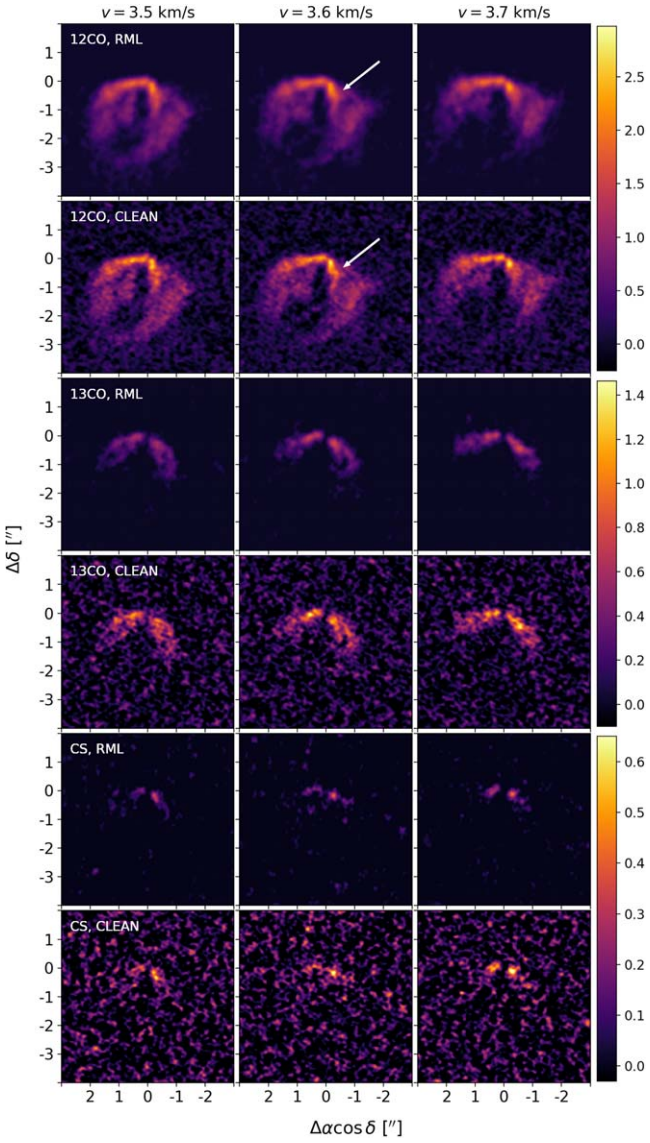

Figure 8. Same as Figure 2, but for SY Cha. Shown are three adjacent channels, centered at $v = 3.6$ km s^{-1} where the ^{12}CO $J = 3$–2 non-Keplerian feature appears most prominently.

is clearly reproduced in the RML images. This feature is consistent with a 1 $M_{\rm Jup}$ planet embedded in the disk (C. Pinte et al. 2025). Like J1615, the CS $J = 7$–6 emission for J1842 is dominated by noise and appears more prominently in the RML images due to the sparsity regularization. However, the structure of the CS $J = 7$–6 emission is not clear in either the RML or CLEAN images at this sensitivity; further spectral averaging could reveal the CS $J = 7$–6 morphology with more clarity.

4.6. LkCa 15

The CLEAN images of LkCa 15 show an NKF in the ^{12}CO $J = 3$–2 emission centered at roughly $v = 6.9$ km s^{-1}. Figure 7 shows a comparison of the RML and CLEAN images for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. All of the RML images in Figure 7 were made with $\lambda_{\rm ent} = 0$, $\lambda_{\rm TSV} = 5 \times 10^{-4}$, and $\lambda_{\rm spa} = 1 \times 10^{-5}$.

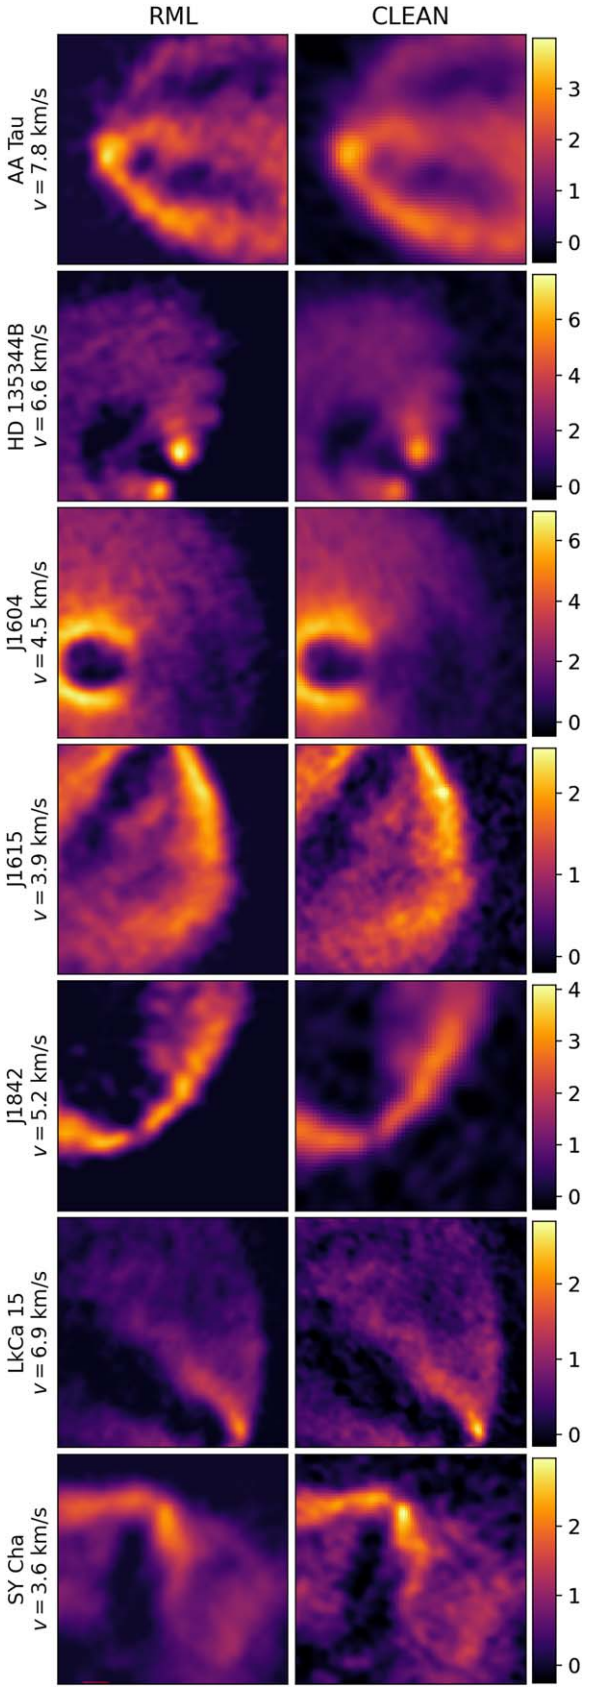

Figure 9. Zoomed-in view of the ^{12}CO $J = 3$–2 NKFs present in the RML (left) and CLEAN (right) images for each source. The color bars are in units of Jy arcsec^{-2}.

The NKF in the ^{12}CO $J = 3$–2 emission is present in both the RML and CLEAN images and appears across all three channels pictured. This feature is consistent with a $2\,M_{\rm Jup}$ planet embedded in the disk (C. Pinte et al. 2025). The RML images reproduce all features seen in the fiducial CLEAN images, though the CLEAN images are slightly sharper. The RML features more uniform background emission due to the sparsity regularization.

4.7. SY Cha

The CLEAN images of SY Cha show an NKF in the ^{12}CO $J = 3$–2 emission, centered at roughly $v = 3.6$ km s^{-1}. Figure 8 shows a comparison of the RML and CLEAN images of SY Cha for the continuum-subtracted ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 data. All of the RML images in Figure 8 were made with $\lambda_{\rm ent} = 0$, $\lambda_{\rm TSV} = 5 \times 10^{-4}$, and $\lambda_{\rm spa} = 1 \times 10^{-5}$.

The emission in the RML images generally appears smoother and more blurred compared to the CLEAN images, but all features seen in the CLEAN images are reproduced, including the NKF in multiple channels. This feature is consistent with a $5\,M_{\rm Jup}$ planet embedded in the disk (C. Pinte et al. 2025). The sparsity regularization heavily suppresses background noise in the RML images compared to the CLEAN images, particularly in the fainter ^{13}CO $J = 3$–2 and CS $J = 7$–6 emission.

5. Discussion

5.1. Resolution and Noise of RML Images

Because RML images do not need to be convolved with a restoring beam (as is common practice in the final step of synthesizing CLEAN images), the spatial resolution of RML images can be difficult to define precisely. Taking the extra step to convolve RML images with a restoring beam is a straightforward way to place a firm resolution limit on the image. While RML images can often achieve a modest degree of super-resolution, the theoretical limit is roughly a factor of 4 compared to the nominal resolution of the interferometer $R_{\rm min} = \lambda / b_{\rm max}$, where $b_{\rm max}$ is the longest baseline in the array. However, given that there is no measured information on the scale of these super-resolved features, this is only possible for data with extremely high SNRs and (u, v) coverage combined with regularizers that are well-matched to the source structure (e.g., R. Narayan & R. Nityananda 1986; M. A. Holdaway 1990; M. Honma et al. 2014). Some RML imaging applications to EHT data have achieved resolutions of 25%–30% of the diffraction limit (K. Akiyama et al. 2017a, 2017b), while others find resolutions of ∼30% of the CLEAN beam (K. Kuramochi et al. 2018). Studies that apply RML imaging techniques to ALMA protoplanetary disk observations typically quote the resolution in relation to the CLEAN beam (which depends on the data weighting and is often significantly larger than the diffraction limit), achieving angular resolutions of up to 1/3 the nominal CLEAN beam (e.g., M. Yamaguchi et al. 2020, 2021, 2024; B. Zawadzki et al. 2023). We thus generate two sets of beam-convolved RML cubes: one with a circular 0.″15 convolution kernel (equivalent to the fiducial CLEAN beam) and one with a circular 0.″05 convolution kernel (1/3 of fiducial CLEAN beam).

A comparison of the fiducial 0.″15 CLEAN, 0.″15 RML, 0.″05 RML, and native RML (no restoring beam) images is shown in Figure 10. We show a single channel ($v = 5.2$ km s^{-1}) of ^{12}CO $J = 3$–2 emission for J1842 for each image product. The 0.″15 RML and CLEAN images show that the RML images reproduce all of the features seen in the CLEAN images, with the only major difference being the background noise suppression in the RML image. We note that a convolution kernel of a given size does not necessarily result in images with an equivalent resolution, as the true resolution of the restored image depends on both the native resolution of the RML image as well as the size of the convolution kernel. The true resolution of the convolved cube can be estimated by adding the native RML image resolution and the convolution kernel size in quadrature. For example, if the native resolution of the RML image of J1824 (rightmost panel, Figure 10) is 0.″05, then convolving the image with a 0.″15 convolution kernel results in a restored image with a resolution of 0.″158. As the resolution of the native RML image cannot easily be determined, may vary spatially, and differs from source to source, we have not attempted to correct for this effect. As a result, the stated resolutions of our beam-convolved RML image cubes are lower limits on the true resolutions of the images.

With a ground truth image for comparison, it would be possible to convolve the image with a variety of beam sizes to determine the resolution that minimizes the normalized rms error. This is commonly used to evaluate the quality of a reconstructed image (e.g., A. A. Chael et al. 2016; K. Akiyama et al. 2017a, 2017b; K. Kuramochi et al. 2018; M. Yamaguchi et al. 2020; B. Zawadzki et al. 2023), but without a "true" image it is not possible to firmly identify an optimal restoring beam size.

Estimating the noise of RML images also poses challenges, as regularization (particularly sparsity regularization) can effectively suppress background noise beyond the inherent thermal noise expected from the instrument. Figure 11 shows the rms noise in a signal-free region of three image products: a native RML, 0.″15 RML, and 0.″15 CLEAN channel of the ^{12}CO $J = 3$–2 emission in AA Tau. The 0.″15 RML and 0.″15 CLEAN images offer the most direct comparison, as the image products are in the same units (Jy beam^{-1}, with identical beam sizes). The rms noise is nearly a factor of 100 smaller in the 0.″15 RML image compared to the 0.″15 CLEAN image (0.043 mJy beam^{-1} and 3.953 mJy beam^{-1}, respectively).

This is because RML and CLEAN images are synthesized with fundamentally different procedures; while the CLEANing process results in a residual map, RML images are subject to regularizer penalties which can favor image qualities like sparsity or smoothness. With sparsity regularization in particular, when the data do not support any significant emission in a region (as in the noisy background of a high-SNR image), the pixel values will tend to zero. This has the overall effect of heavily suppressing noise in signal-free regions of the image, far beyond expected values from the thermal noise of the instrument. This does not mean that the overall noise in the RML images is truly orders of magnitude better than the corresponding CLEAN products. Rather, the noise across an RML image may not be spatially uniform, as sparsity regularization only suppresses pixel values in regions where there is not likely to be any emission at all.

5.2. Emission Surfaces and Temperatures

We compared the nonparametric emission surfaces and radial temperature profiles, presented in full in M. Galloway-Sprietsma et al. (2025), between the CLEAN cubes and the RML cubes. To

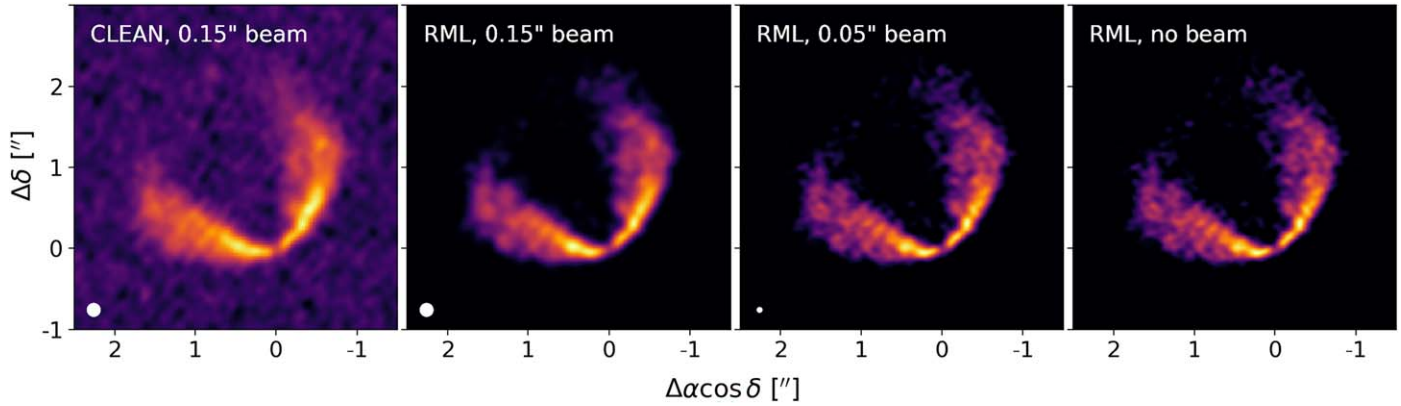

Figure 10. Left to right: fiducial 0.″15 CLEAN, 0.″15 RML, 0.″05 RML, and native RML (no restoring beam) images. We show each image product for the continuum-subtracted ^{12}CO $J = 3$–2 emission in J1842 at $v = 5.2$ km s^{-1}, with each panel plotted on an independent color scale. While beam convolution is not necessary in the RML imaging process, it can be useful to enable more direct comparisons between CLEAN and RML, as well as for constraining the resolution of the RML images. Beam-convolved images are also necessary for compatibility with some analysis software (e.g., tasks that require information about the restoring beam, which does not exist in the native RML image).

compare the two, we used RML cubes of the same spatial and spectral resolution as those used for the CLEAN cube analysis (0.″15, 100 m s^{-1}). We use the non-continuum-subtracted cubes so as to not underestimate the temperature.

We follow identical procedures to those outlined in Section 3 of M. Galloway-Sprietsma et al. (2025) to obtain the RML emission surfaces. The procedure is summarized as follows. We utilize the open-source Python package `disksurf`, which obtains emission surfaces via the method first demonstrated in C. Pinte et al. (2018). To trace the emission surface, `disksurf` locates intensity maxima along a vertical line intersecting isovelocity contours. For disks of moderate inclination, the front and back surfaces will be well-separated, allowing us to distinguish between the front and back surfaces, and to trace the $r - z$ location of emission. As in M. Galloway-Sprietsma et al. (2025), we do this process over five iterations, each time increasing the minimum SNR from 1.0 to 5.0 for ^{12}CO $J = 3$–2 and 1.0 to 3.5 for ^{13}CO $J = 3$–2. Minimal masking is applied to the surfaces so that we can assess the differences between CLEAN and RML. This allows us to obtain the "raw" surfaces, comprising $r - z$ points. We completed this procedure with the RML cubes of AA Tau, J1615, J1842, LkCa 15, and SY Cha for both ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2 emission.

Figure 12 shows the results for J1615 and SY Cha. The top two panels show ^{12}CO $J = 3$–2 emission surfaces, and the bottom two show ^{13}CO $J = 3$–2 surfaces. The raw $r - z$ points are shown in the background in gray for the CLEAN cubes, and in aqua blue (^{12}CO $J = 3$–2) and red (^{13}CO $J = 3$–2) for the RML cubes. The larger points show the raw surfaces binned by a quarter of the beam size for visual clarity. The dashed vertical lines show the maximum radial extent of the retrieved surfaces. In general, the emission surfaces are in agreement, which is also what we find for the nonpictured disks (AA Tau, J1842, and LkCa 15). The J1615 emission surfaces both have the same characteristic tapered power-law shape. The morphology of the SY Cha ^{12}CO $J = 3$–2 RML emission surface has the biggest difference among the five disks tested; the surfaces begin to diverge at ∼300 au. Past 300 au, the CLEAN emission surface continues to rise, whereas the RML surface begins to drop down. The reasons for these morphological differences are unclear. Eventually, both surfaces drop at ∼400 au.

Additionally, the RML emission surfaces tend to extend further out in both ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2. This can be seen in nearly all of the surfaces shown in Figure 12. One of the most dramatic differences is in SY Cha ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2 emission, whose RML surface extends over 200 au further than that found using the CLEAN cube for both molecules. The effect is less pronounced in J1615, but there is still a ∼100 au size difference between the ^{12}CO $J = 3$–2 emission surfaces. This is likely due to the decreased rms noise present in the RML cube versus the CLEAN cube. With lower noise, the `disksurf` surface-finding process will be able to identify more points that were previously cut off by the 5 SNR noise limit.

In addition to the emission surfaces, we compare the radial temperature profiles between the CLEAN and RML cubes. We obtain the temperature as a function of the radius using the $r - z$ points extracted by `disksurf`. For further details, see Section 4 of M. Galloway-Sprietsma et al. (2025). Figure 13 shows a comparison of the brightness temperatures between the CLEAN cubes and the RML cubes for ^{12}CO $J = 3$–2 (top panels) and ^{13}CO $J = 3$–2 (bottom panels). For visual clarity, we have binned the data by a quarter of the beam size for each corresponding cube. As noted in the previous paragraph, the maximum radial extent is larger for the RML cubes. We find that the temperatures derived from the RML cubes are systematically lower at each radius; this can be seen in the offset between the binned points. For the ^{12}CO $J = 3$–2 emission, the difference, on average, is only ∼3 K, but in some places, it becomes as large as 13 K, such as in the inner radial regions of SY Cha. For J1615 in ^{13}CO $J = 3$–2, the temperature difference, on average, is only 1.3 K. The ^{13}CO $J = 3$–2 emission in SY Cha exhibits the largest temperature difference between the CLEAN and RML cubes, with the average offset being 8 K.

The causes of the temperature offset between the CLEAN and RML cubes may, in part, be due to resolution differences. To test this, we derived the temperatures using RML cubes of SY Cha and AA Tau convolved with multiple beam sizes (0.″5, 0.″15, and 0.″30). These sources were selected because the native RML images of SY Cha seem to have a slightly worse resolution than the fiducial CLEAN cubes, while the native RML images of AA Tau appear to be super-resolved. The SY

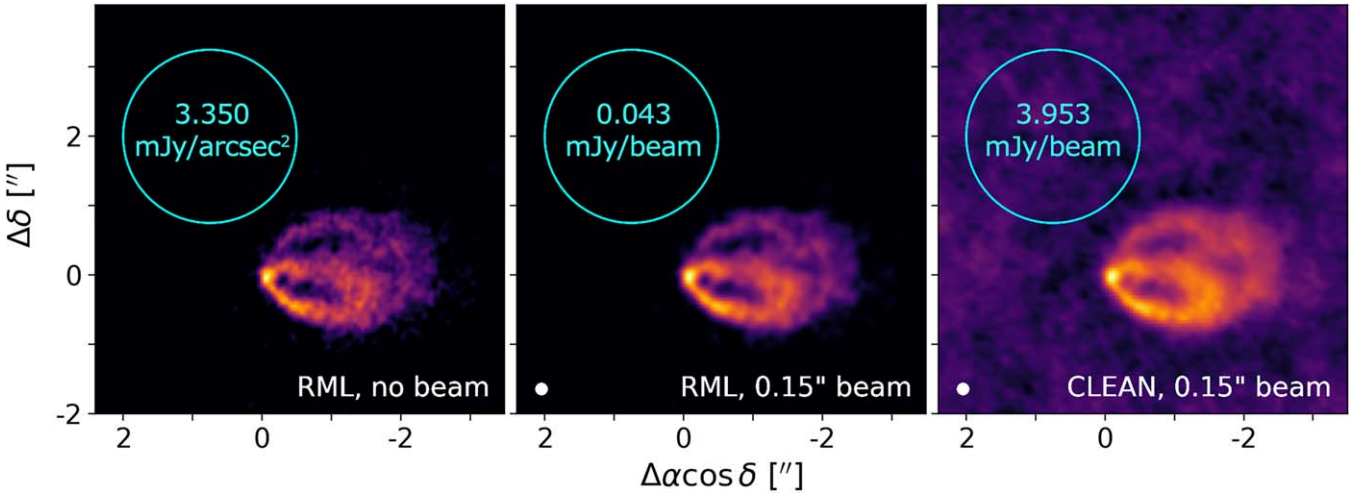

Figure 11. Three different image products showing continuum-subtracted ^{12}CO $J = 3$–2 emission in AA Tau at $v = 7.8$ km s^{-1}. We estimate the noise by taking the rms of a signal-free region, shown with the cyan circles. Before beam convolution, the native RML images are in units of Jy arcsec^{-2} (left). To enable a more direct comparison of the noise properties between RML and CLEAN images, we calculate the rms in the RML image convolved with a circular 0.″15 beam (center) along with the fiducial CLEAN image, which also has a circular 0.″15 beam (right). The rms in the 0.″15 RML image is nearly 2 orders of magnitude lower than the rms in the 0.″15 fiducial CLEAN image (0.043 mJy beam^{-1} and 3.953 mJy beam^{-1}, respectively), demonstrating that regularization can suppress the background noise far below what is expected from instrumental thermal noise.

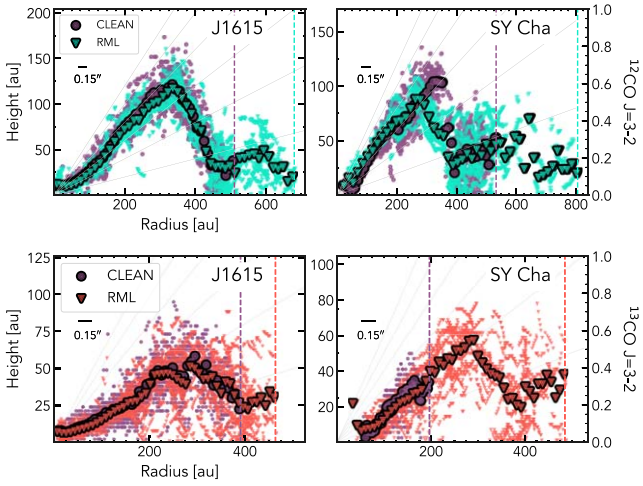

Figure 12. Emission surfaces for ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2 for J1615 and SY Cha found using the CLEAN cubes and the RML cubes. Raw $r - z$ points are shown in purple for the CLEAN data and in aqua blue and red for the RML data. The larger points show the binned surfaces. The horizontal dashed lines show the maximum surface radius, $r_{\rm max}$. The beam size of 0.″15 is shown in the upper left corner.

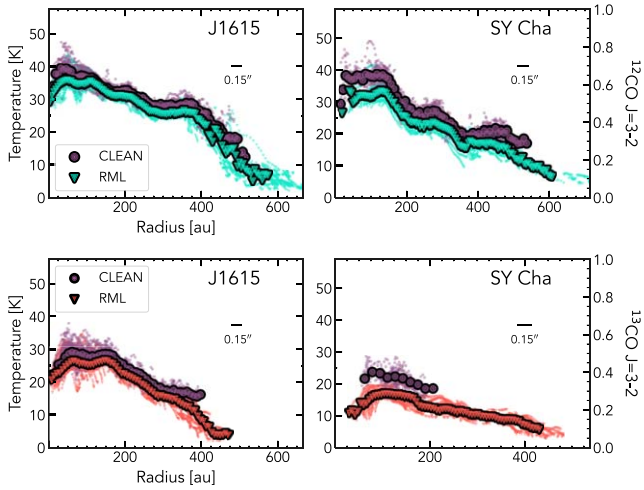

Figure 13. Radial temperature profiles, found using the `disksurf` surface points, for ^{12}CO $J = 3$–2 and ^{13}CO $J = 3$–2 for J1615 and SY Cha using the CLEAN and RML cubes. The larger points show the temperatures binned by a quarter of the beam size. The beam size of 0.″15 is shown in the upper left corner.

Cha temperatures extracted from the 0.″05 and 0.″15 RML images were indistinguishable, while the 0.″30 RML images yielded a slightly lower temperature. This behavior is observed to a larger extent with the CLEAN images, where the 0.″30 images result in significantly lower derived temperatures compared to the fiducial 0.″15 images (and also significantly lower than temperatures derived from any of the 0.″05, 0.″15, or 0.″30 RML images). The AA Tau temperatures from the 0.″5, 0.″15, and 0.″30 images were consistent with each other, with only minor differences in the 0.″05 temperature profile. In particular, the derived temperature profile for the 0.″05 cube was not consistently hotter than the 0.″15 and 0.″30 temperatures but was hotter at some radii. This suggests that disk temperature profiles derived from RML images may be more robust to resolution differences than CLEAN images, though further analysis is needed to characterize this behavior in detail.

Another possible explanation for the temperature offset between the CLEAN and RML cubes is the lower rms of the RML images. It may be that the additional emission surface points found with the RML cubes belong to colder regions of the disk. These points may be cut out by the SNR clip applied when using `disksurf` on the CLEAN cubes, but with a lower-noise RML cube, they could remain, thus causing the offset. Another explanation could be the imaging process itself. The final intensities, and thus temperatures, are dependent on the imaging process. Fundamental differences between CLEAN and RML could lead to slight offsets in the intensities. Further exploration of this is needed. However, despite these offsets, we find consistent radial profile morphology, and most

of the temperature differences are within the intrinsic temperature scatter of ∼5 K.

5.3. Limitations and Future Work

In this work, we have presented RML images of seven exoALMA protoplanetary disks, with the primary goal of detecting protoplanets at high confidence. However, RML imaging can have other benefits beyond validating CLEAN images, namely, the potential for improved resolution in comparison to the CLEAN beam (see Section 5.1). We do not achieve a clear improvement in resolution in all the images presented here. This is likely due to the fact that we have treated all disks uniformly, i.e., tuning hyperparameters using strictly CV. This could be improved by repeating the CV procedure with different visibility partitioning or by fine-tuning hyperparameters by hand following the CV process. A more individual RML exploration of each disk could result in a higher degree of super-resolution in the images, potentially reaching resolutions several times better than the fiducial CLEAN images and enabling the potential to reveal new substructure not resolved in the CLEAN images. This detailed and source-specific procedure would require large amounts of human input and computational power and thus is not practical for a first analysis of a large sample like the exoALMA data. However, future work should include detailed RML imaging of the full exoALMA data set (and other protoplanetary disk data of comparable quality) with the explicit goal of obtaining super-resolved images and searching for previously unseen structures, particularly in disks for which we currently do not see evidence of embedded protoplanets. While the CLEAN images would not necessarily be able to support the detection of such features, comparison with other RML algorithms could provide a path toward confidently identifying features not seen with CLEAN imaging.

This would also require imaging the RML images with smaller pixel sizes to ensure that the model parameterization enables the highest potential degree of super-resolution. Given that the highest resolution data from the interferometer corresponds to resolutions of about 0.″06, the pixel size of 0.″0125 used in this study may not sufficiently sample the image for the best possible degree of super-resolution relative to the nominal resolution of the interferometer (but does not affect the results presented here). Pushing to a smaller pixel size would further increase the computational demands of the imaging process and would require a new suite of hydrodynamical simulations for comparison. It would also be worthwhile to obtain RML images at the highest spectral resolution possible from the data (i.e., 28 m s^{-1}), as careful RML imaging may be able to mitigate the lower SNR without the spectral averaging, which would further increase the computational load. Nevertheless, this may be the best path forward for identifying the smallest protoplanets with the currently available data.

Additionally, future work should explore the impact of normalization schemes for different regularizers in greater depth. Here, we find that optimal hyperparameter values do not vary significantly across the cube, observing stable performance across channels even without explicitly normalizing the sparsity and TSV regularizers. However, a more detailed investigation of regularizer normalization for ALMA protoplanetary disk data, such as testing the various normalization factors presented in Event Horizon Telescope Collaboration et al. (2019), could offer further improvements. In particular, incorporating normalization factors to account for flux variations across velocity space may enhance the image quality for fainter emission, such as the ^{13}CO $J = 3$–2, and CS $J = 7$–6 data, or in the wings of the line emission, where small variations in the regularizer strength could have a stronger effect.

6. Conclusions

We have presented RML images for seven exoALMA targets. Our main conclusions are summarized below.

1. The RML methods independently and consistently reproduce the features seen in the fiducial CLEAN images, including all NKFs. The agreement between the two sets of independently synthesized image products suggests that these features are real and strengthens findings related to the interpretation of these features (such as those presented in C. Pinte et al. 2025).
2. Additionally, RML methods generally reproduce the azimuthally averaged emission surfaces and temperature profiles found with CLEAN cubes. The emission surface morphology and features are consistent, but the surfaces of the RML cubes extend further radially, sometimes by 200 au. We find that the temperature profiles measured with the RML cubes are systematically lower by ∼5 K, which varies radially and on a per-disk basis. This is likely due to the artificially high SNR of the RML cubes stemming from sparsity regularization, but further exploration is needed to fully characterize this effect.
3. Using CV for hyperparameter tuning, we found that the set of hyperparameters that minimized the CV score (corresponding to a model with the highest predictive power) remained stable from channel to channel in a given cube, having only minor variations that did not significantly change the resulting image products.
4. Furthermore, we find that optimal hyperparameters remain stable for all of the ^{12}CO $J = 3$–2, ^{13}CO $J = 3$–2, and CS $J = 7$–6 images for a given source, suggesting that optimal hyperparameter values depend more on observational parameters like interferometer baselines and integration time rather than the specific emission morphologies.

For each target, the following data products will be publicly available:

1. Full ^{12}CO $J = 3$–2 RML cubes (both continuum-subtracted and non-continuum-subtracted).
2. Full ^{13}CO $J = 3$–2 RML cubes (both continuum-subtracted and non-continuum-subtracted).
3. Full CS $J = 7$–6 RML cubes (both continuum-subtracted and non-continuum-subtracted).

Different versions will be released with the native RML resolution (no beam convolution), convolution with a 0.″05 restoring beam, and convolution with a 0.″15 restoring beam. We have also released sample imaging scripts, available directly on the `MPoL` GitHub.[29]

Acknowledgments

We thank the anonymous reviewers for the comments, which greatly improved this manuscript. Support for B.Z. was

[29] https://github.com/MPoL-dev/examples

provided by The Brinson Foundation. Computations for this research were performed in part on the Pennsylvania State University's Institute for Computational and Data Sciences' Roar supercomputer. We acknowledge the U.S. National Science Foundation—Major Research Instrumentation Program, Deep Learning for Statistics, Astrophysics, Geoscience, Engineering, Meteorology and Atmospheric Science, Physical Sciences and Psychology (DL-SAGEMAPP) at the Institute for Computational and Data Sciences (ICDS) at the Pennsylvania State University. This work was performed in part on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government and from the Victorian Higher Education State Investment Fund (VHESIF) provided by the Victorian Government.

This Letter makes use of the following ALMA data: ADS/JAO.ALMA#2021.1.01123.L. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada), MOST 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. We thank the North American ALMA Science Center (NAASC) for their generous support including providing computing facilities and financial support for student attendance at workshops and publications.

J.B. acknowledges support from NASA XRP grant No. 80NSSC23K1312. M.B., D.F., and J.S. have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (PROTOPLANETS, grant agreement No. 101002188). Computations by JS have been performed on the "Mesocentre SIGAMM" machine, hosted by Observatoire de la Cote d'Azur. P.C. acknowledges support by the Italian Ministero dell'Istruzione, Università e Ricerca through the grant Progetti Premiali 2012—iALMA (CUP C52I13000140001) and by the ANID BASAL project FB210003. S.F. is funded by the European Union (ERC, UNVEIL, 101076613) and acknowledges financial contribution from PRIN-MUR 2022YP5ACE. M.F. is supported by a Grant-in-Aid from the Japan Society for the Promotion of Science (KAKENHI: No. JP22H01274). C.H. acknowledges support from NSF AAG grant No. 2407679. T.H. is supported by an Australian Government Research Training Program (RTP) Scholarship. J.D.I. acknowledges support from an STFC Ernest Rutherford Fellowship (ST/W004119/1) and a University Academic Fellowship from the University of Leeds. A.I. acknowledges support from the National Aeronautics and Space Administration under grant No. 80NSSC18K0828. Support for AFI was provided by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51532.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. C.L. has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No. 823823 (DUSTBUSTERS) and by the UK Science and Technology Research Council (STFC) via the consolidated grant ST/W000997/1. C.P. acknowledges Australian Research Council funding via FT170100040, DP18010423, DP220103767, and DP240103290. G.R. acknowledges funding from the Fondazione Cariplo, grant No. 2022-1217, and the European Research Council (ERC) under the European Union's Horizon Europe Research & Innovation Programme under grant agreement No. 101039651 (DiscEvol). H.-W.Y. acknowledges support from the National Science and Technology Council (NSTC) in Taiwan through grant NSTC 113-2112-M-001-035- and from the Academia Sinica Career Development Award (AS-CDA-111-M03). G.W.F. acknowledges support from the European Research Council (ERC) under the European Union Horizon 2020 research and innovation program (grant agreement No. 815559 (MHDiscs)). G.W.F. was granted access to the HPC resources of IDRIS under the allocation A0120402231 made by GENCI. A.J.W. has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 101104656. T.C.Y. acknowledges support by Grant-in-Aid for JSPS Fellows JP23KJ1008. Views and opinions expressed by ERC-funded scientists are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Software: Astropy (Astropy Collaboration et al. 2013, 2018), PyTorch (A. Paszke et al. 2019), MPoL (I. Czekala et al. 2021; B. Zawadzki et al. 2023; I. Czekala et al. 2025), CASA (J. P. McMullin et al. 2007), disksurf (R. Teague et al. 2021).

ORCID iDs

Brianna Zawadzki https://orcid.org/0000-0001-9319-1296
Ian Czekala https://orcid.org/0000-0002-1483-8811
Maria Galloway-Sprietsma https://orcid.org/0000-0002-5503-5476
Jaehan Bae https://orcid.org/0000-0001-7258-770X
Marcelo Barraza-Alfaro https://orcid.org/0000-0001-6378-7873
Myriam Benisty https://orcid.org/0000-0002-7695-7605
Gianni Cataldi https://orcid.org/0000-0002-2700-9676
Pietro Curone https://orcid.org/0000-0003-2045-2154
Stefano Facchini https://orcid.org/0000-0003-4689-2684
Daniele Fasano https://orcid.org/0000-0003-4679-4072
Mario Flock https://orcid.org/0000-0002-9298-3029
Misato Fukagawa https://orcid.org/0000-0003-1117-9213
Himanshi Garg https://orcid.org/0000-0002-5910-4598
Cassandra Hall https://orcid.org/0000-0002-8138-0425
Thomas Hilder https://orcid.org/0000-0001-7641-5235
Jane Huang https://orcid.org/0000-0001-6947-6072
John D. Ilee https://orcid.org/0000-0003-1008-1142
Andrea Isella https://orcid.org/0000-0001-8061-2207
Andrés F. Izquierdo https://orcid.org/0000-0001-8446-3026
Kazuhiro Kanagawa https://orcid.org/0000-0001-7235-2417
Geoffroy Lesur https://orcid.org/0000-0002-8896-9435
Cristiano Longarini https://orcid.org/0000-0003-4663-0318
Ryan A. Loomis https://orcid.org/0000-0002-8932-1219
Ryuta Orihara https://orcid.org/0000-0003-4039-8933
Christophe Pinte https://orcid.org/0000-0001-5907-5179
Daniel J. Price https://orcid.org/0000-0002-4716-4235
Giovanni Rosotti https://orcid.org/0000-0003-4853-5736
Jochen Stadler https://orcid.org/0000-0002-0491-143X
Richard Teague https://orcid.org/0000-0003-1534-5186
Hsi-Wei Yen https://orcid.org/0000-0003-1412-893X

Gaylor Wafflard-Fernandez https://orcid.org/0000-0002-3468-9577
David J. Wilner https://orcid.org/0000-0003-1526-7587
Andrew J. Winter https://orcid.org/0000-0002-7501-9801
Lisa Wölfer https://orcid.org/0000-0002-7212-2416
Tomohiro C. Yoshida https://orcid.org/0000-0001-8002-8473

References

Aizawa, M., Suto, Y., Oya, Y., Ikeda, S., & Nakazato, T. 2020, ApJ, 899, 55
Akiyama, K., Ikeda, S., Pleau, M., et al. 2017b, AJ, 153, 159
Akiyama, K., Kuramochi, K., Ikeda, S., et al. 2017a, ApJ, 838, 1
Andrews, S. M., Huang, J., Pérez, L. M., et al. 2018, ApJL, 869, L41
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
Bishop, C. M. 2006, in Pattern Recognition and Machine Learning, ed. M. Jordan, J. Kleinberg, & S. Bernhard (New York: Springer)
Breiman, L., & Spector, P. 1992, Int. Stat. Rev., 60, 291
Cárcamo, M., Muñoz, N., Rannou, F., et al., 2019 GPUVMEM: Maximum Entropy Method (MEM) GPU Algorithm for Radio Astronomical Image Synthesis, Astrophysics Source Code Library, ascl:1906.014
Cárcamo, M., Román, P. E., Casassus, S., Moral, V., & Rannou, F. R. 2018, A&C, 22, 16
Casassus, S., Cárcamo, M., Hales, A., Weber, P., & Dent, B. 2022, ApJL, 933, L4
Casassus, S., Christiaens, V., Cárcamo, M., et al. 2021, MNRAS, 507, 3789
Casassus, S., & Pérez, S. 2019, ApJL, 883, L41
CASA Team, Bean, B., Bhatnagar, S., et al. 2022, PASP, 134, 114501
Chael, A. A., Johnson, M. D., Bouman, K. L., et al. 2018, ApJ, 857, 23
Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, ApJ, 829, 11
Cornwell, T. J., & Evans, K. F. 1985, A&A, 143, 77
Czekala, I., Jennings, J., Zawadzki, B., et al. 2023, MPoL-dev/MPoL: v0.2.0 Release, Zenodo, doi: 10.5281/zenodo.10107615
Czekala, I., Jennings, J., Zawadzki, B., et al. 2025, arXiv:2502.00100
Czekala, I., Zawadzki, B., Loomis, R., et al. 2021, MPoL-dev/MPoL: v0.1.1 Release, Zenodo, doi: 10.5281/zenodo.4939048
Disk Dynamics Collaboration, Armitage, P. J., & Bae, J. 2020, arXiv:2009.04345
Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, ApJL, 875, L4
Galloway-Sprietsma, M., Bae, J., Izquierdo, A., et al. 2025, ApJL, 984, L10
Gull, S. F., & Daniell, G. J. 1978, Natur, 272, 686
Högbom, J. A. 1974, A&AS, 15, 417
Högbom, J. A. 1979, in IAU Colloq. 49: Image Formation from Coherence Functions in Astronomy, ed. C. van Schooneveld (Cambridge: Cambridge Univ. Press), 237
Holdaway, M. A. 1990, PhD thesis, Brandeis Univ., Waltham, MA
Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, PASJ, 66, 95
Kohavi, R. 1995, in Proc. 14th Int. Joint Conf. on Artificial Intelligence—Vol. 2, IJCAI'95 (San Francisco, CA: Morgan Kaufmann Publishers Inc), 113
Kuramochi, K., Akiyama, K., Ikeda, S., et al. 2018, ApJ, 858, 56
Loomis, R., Facchini, S., Benisty, M., et al. 2025, ApJL, 984, L7
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
Molinaro, A. M., Simon, R., & Pfeiffer, R. M. 2005, Bioin, 21, 3301
Narayan, R., & Nityananda, R. 1986, ARA&A, 24, 127
Paszke, A., Gross, S., Massa, F., et al. 2019, in Proc. of the 33rd Int. Conf. on Neural Information Processing Systems, PyTorch: an imperative style, high-performance deep learning library, ed. H. Wallach et al. (Red Hook, NY: Curran Associates, Inc.), 8026
Pérez, S., Casassus, S., Baruteau, C., et al. 2019, AJ, 158, 15
Pérez, S., Casassus, S., Hales, A., et al. 2020, ApJL, 889, L24
Pinte, C., Ilee, J., Huang, J., et al. 2025, ApJL, 984, L15
Pinte, C., Ménard, F., Duchêne, G., et al. 2018, A&A, 609, A47
Rau, U., & Cornwell, T. J. 2011, A&A, 532, A71
Schwarz, U. J. 1978, A&A, 65, 345
Teague, R., Benisty, M., Facchini, S., et al. 2025, ApJL, 984, L6
Teague, R., Law, C., Huang, J., & Meng, F. 2021, JOSS, 6, 3827
Tibshirani, R. 1996, Journal of the Royal Statistical Society. Series B (Methodological), 58, 267
Yamaguchi, M., Akiyama, K., Tsukagoshi, T., et al. 2020, ApJ, 895, 84
Yamaguchi, M., Muto, T., Tsukagoshi, T., et al. 2024, PASJ, 76, 437
Yamaguchi, M., Tsukagoshi, T., Muto, T., et al. 2021, ApJ, 923, 121
Zawadzki, B., Czekala, I., Loomis, R. A., et al. 2023, PASP, 135, 064503