European Organisation for Astronomical Research in the Southern Hemisphere

Atacama Large Millimeter/submillimeter Array

ESO ALMA Support Centre

Internal ALMA Development Study

ALMA 2030: Development Study on Bayesian Adaptive Interferometric Image Reconstruction Methods

Final REPORT

Fabrizia Guglielmetti (ESO)

September 2024

Abridged

The Atacama Large Millimeter/submillimeter Array (ALMA) will achieve unprecedented amount of deep, high-resolution observations with its upcoming upgrades. The planned advancements will significantly enhance ALMA's spectral bandwidth, accelerate spectral scanning, and improve sensitivity across all observations. While these upgrades promise new scientific capabilities, they also present significant operational challenges, particularly in data reduction, quality assurance, and data management.

Image reconstruction is inherently complex due to factors such as sparse sampling, varying instrumental responses, diverse celestial source morphologies, a range of astrophysical processes and pervasive noise. Data volume handling of the imaging products is also a challenge. Currently, ALMA's interferometry imaging products are streamlined using mitigation techniques to reduce computational costs, which facilitates quicker delivery of calibration and imaging products to the Principal Investigators. However, this approach has led to the ALMA Science Archive missing many images resulting from the data reduction pipeline's mitigation efforts.

ALMA currently generates 1 TB of scientific data daily, a rate expected to increase by at least an order of magnitude over the next decade. Upgrades to receivers and the correlator will enhance sensitivity and efficiency, resulting in single-field and mosaic cubes that are at least two orders of magnitude larger than the current gigabyte-sized cubes. While these advancements will significantly boost observational capabilities, they will also heighten computational demands for data reduction and quality assurance.

Exploration and testing of alternative image processing techniques are already essential. Innovative solutions are necessary to ensure a complete archive of imaging products and facilitate effective data mining. Advanced algorithms are required for reducing processing time and for managing the big data regime, as the number of observed spectral lines doubles and resolution improves. The employed algorithms have to deliver robust and reliable results to minimize the need for human intervention.

To address these challenges, we propose exploring artificial intelligence methodologies for imaging to enhance ALMA's scientific impact. This includes tackling issues such as thresholding, continuum subtraction, detection of extended emissions, separation of point-like sources from diffuse emissions, weak signal detection, mosaic analysis, and improving processing speeds. Imaging algorithms have been tailored using astro-statistics and astro-informatics techniques, i.e. `RESOLVE` and `DeepFocus`. With the aim in mind to enhance supervised ML techniques, a new ALMA simulator `ALMASim` has been developed. `ALMASim` is aided by an empirical noise simulator `NOISEMPIRE` to address realistic noise properties in ALMA simulated data. All software utilized in this study is openly accessible to the scientific community, giving them extremely high legacy value. Looking ahead, the need to further develop, test and optimize artificial intelligence techniques is even more crucial. These efforts should prioritize comparing the accuracy, precision, speed, and data-handling capacity of each software to establish performance benchmarks. Such comparative analysis will pave the way for new opportunities, including the development of science platforms for the user community and software optimized for real-time image analysis in support of various observatories.

Executive Summary

This study presents an initial exploration of applying artificial intelligence methodologies to ALMA data for imaging purposes. The advanced imaging techniques are designed to handle large data volumes efficiently, reducing the need for human intervention while enhancing data processing capabilities. The ultimate goal is to equip the scientific community with sophisticated tools that produce high-quality, comprehensive imaging products. By optimizing these tools to manage data volume requirements and minimize computational costs, our approach will unlock new opportunities for data mining and revolutionize observatory operations. Additionally, new user-friendly solutions can be developed to support the next generation of users, including non-experts.

We utilize two distinct software tools to analyze ALMA data, `RESOLVE` [1] and `DeepFocus` [2], addressing

the challenges outlined in the ALMA2030 development roadmap [3, 4]. The techniques differ in nature: one is grounded in astro-statistics, while the other focuses on astro-informatics, respectively. Both methods have demonstrated key strengths and possess essential capabilities crucial for the Big Data era. A documentation of the developed software is accessible on the GitHub [1] [5] and GitLab [2] [6] Integration Platforms.

Using real and simulated data sets, we investigate these techniques to tackle the synthesis imaging challenges in view of the ALMA2030 era.

For real ALMA data, we leverage Science Verification (SV) datasets[3], the DSHARP [7] and ALCHEMI [8] Large Programs, as well as other publicly available datasets. The SV process ensures data quality for scientific analysis, preparing the products for thorough quality assessment. Our analysis includes iconic datasets such as HLTau and BR1202. The DSHARP and ALCHEMI Large Programs provide an excellent foundation for evaluating `RESOLVE`'s performance in detecting extended emissions and applying a probabilistic ALMA GOUS imaging approach to combining interferometric data with single-dish data. Furthermore, we test the algorithm for its ability to detect the Sunyaev-Zel'dovich effect and analyze data related to the Circum Galactic Medium. Additionally, the `DeepFocus` algorithm and its companion software `ALMASim`[4] [9] are optimized for the serendipitous detection of faint QSOs within a known sample of ALMA archival data.

For ALMA simulated data, we have several options: `ALMASim` [9], `NIFTy`[5] [10] and `CASA` [11]. The latter is the software package commonly used to calibrate, image and simulate ALMA data. The simulated data are employed by the available software and comparisons with `tCLEAN` [12] are performed.

`RESOLVE` is a robust algorithm and founded on a principled method. It outperforms current imaging techniques for the detection of diffuse emission. Complex structures in the celestial signal and point-like sources are well detected. Super-resolution imaging is achieved. The required data for this method are the ALMA observed calibrated visibilities. The input variables, as celestial signal and power spectrum, are initialized and estimated by the data during the optimization to the most probable image configuration. The reconstructed images provide for a reliable solution with no need of extra human intervention. `RESOLVE` is applied to ALMA continuum images and cubes. The application of the `RESOLVE` algorithm for aggregate continuum data, spectral cubes, and mosaicking demonstrates significant potential. However, further advancements are essential to unlock its full capabilities. While `RESOLVE` reconstructions are computationally intensive, considerable progress has been achieved, particularly with the development of `fast-resolve` [13]. Similar optimization efforts are ongoing for other software, such as `CASA tCLEAN` [14]. Future enhancements could include integrating ML methods with RESOLVE, offering opportunities to accelerate computations further and enhance performance. These developments aim to address the balance between computational efficiency and the algorithm's robust imaging quality. Moreover, although `RESOLVE` is computationally expensive (comparable to a `tCLEAN` execution), the algorithm delivers in addition to the deconvolved image other informative products as uncertainty map, power spectrum, and its uncertainty. The uncertainty map is commonly the most computational expensive. In addition, self-calibration and full polarization imaging have been developed within the IFT framework [15–18]. This software has the potential to lay the groundwork for a fully automated pipeline and significantly enhance scientific research.

`DeepFocus` has demonstrated exceptional image fidelity and computational efficiency in reconstructing ALMA data cubes. The technique processes ALMA dirty cubes by learning celestial sources, noise patterns, and the instrumental point spread function directly from the input data. This learning occurs across both spatial and frequency domains, ensuring that information from each channel is propagated throughout the frequency

---

[1] https://github.com/MicheleDelliVeneri/DeepFocus

[2] https://gitlab.mpcdf.mpg.de/ift/resolve

[3] https://almascience.eso.org/alma-data/science-verification

[4] https://github.com/MicheleDelliVeneri/ALMASim

[5] https://github.com/NIFTy-PPL

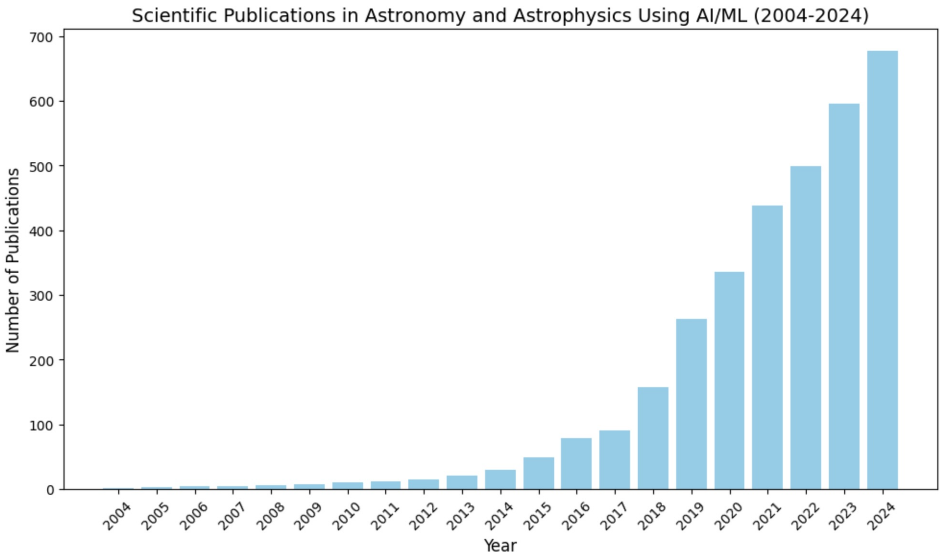

Figure 1: The application of artificial intelligence and ML in astronomy and astrophysics continues to grow rapidly. The volume of publications focusing on these techniques has surged in response to the increasing availability of large datasets from observatories like the Vera Rubin Observatory and space missions like JWST. This increase reflects the shift toward using AI/ML to handle massive datasets that are impractical for human analysis (created on Sep 11, 2024)

spectrum. By effectively leveraging spatial and frequency data, this approach enables extreme data compression. `DeepFocus` is also highly adaptable, thanks to its integration with the `ALMASim` package, which supports training, testing, and validation on a wide range of celestial signals. The coupling of `ALMASim` with the additional tool `NOISEMPIRE`[6] [19] introduces a more accurate representation of noise distributions typical of ALMA data. It separates and categorizes various noise components from archived ALMA images, facilitating the identification of decorrelation patterns caused by noisy antennas or other issues. In addition to enable simulation-based inference, a future potential for `NOISEMPIRE` lies in its application to alarm systems, aimed at improving observatory throughput and enhancing data quality.
Astro-informatics holds significant potential for revolutionizing data management in science archives and operational processes. Supervised ML has the potential to significantly enhance ALMA's operations by supporting real-time image analysis for astronomers on duty, as well as aiding data mining efforts and support the Principal Investigators. This technology can streamline the analysis process, providing quick insights and improving overall data accessibility. Furthermore, automated cataloging of data from each ALMA Cycle would simplify data access, making the archive more efficient.

Based on the current investigations, `RESOLVE` emerges as the optimal algorithm for robust detection of diffuse emission and faint sources. Meanwhile, `DeepFocus`, integrated within `ALMASim`, is well-positioned to address the performance challenges introduced by the Wideband Sensitivity Upgrade (ALMA2030) [3]. As ALMA transitions into the Big Data era, advancements in interdisciplinary domains are essential to address challenges in image reconstruction, classification, and simulation-based inference (Fig. 1). With the planned upgrades, ALMA image analysis increasingly demands adaptive algorithms capable of uncovering patterns in data, learning from new inputs, and maximizing the extraction of meaningful information.

---

[6] https://github.com/Ibaronch/NOISEMPIRE

Contents

Acronyms

ACF	Autocorrelation Function
AI	Artificial Intelligence
ALCHEMI	ALMA Comprehensive High-resolution Extragalactic Molecular Inventory [8]
ALMA	Atacama Large Millimeter/submillimeter Array
ALMASim	ALMA Simulator within the supervised ML contest [9]
AGN	Active Galactic Nucleus
ARC	ALMA Regional Centre
BlobsFinder	detect and analyze structures, commonly used in astroinformatics
BRAIN	Bayesian Reconstruction with Adaptive Image Notion
C2PAP	Computational Center for Particle and Astrophysics
CAE-VAE	Convolutional Autoencoder - Variational Autoencoder, combining CNN and probabilistic models
CASA	Common Astronomy Software Applications [20]
CMB	Cosmic Microwave Background
CNN	Convolutional Neural Network, a Deep Learning NN
CPU	Central Processing Unit
CRF	Correlation Random Field
Dask	open-source parallel computing library in Python
DeepFocus	A meta-learning pipeline powered by deep learning techniques [2]
DeepGRU	GRU in Deep Learning, type of RNN
DenseNet	Dense Convolutional Network with high parameter efficiency
DL	Deep Learning, ML capable to train large NN. It is inspired by the structure and functioning of the human brain
DRM	Data Reduction Manager
DSHARP	Disk Substructures at High Angular Resolution Project [7]
DSFG	Dusty Star Forming Galaxy
EA	East Asia
EASC	European ALMA Support Centre
EB	Execution Block
ECML-PKDD	European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases
EfficientNets	CNN architecture with family of models focused on accuracy optimization
EI	Expected Improvement, common acquisition function in Bayesian Optimization
ESAC	European Science Advisory Committee
ESO	European Southern Observatory
EU	Europe
F2F	Face to Face
FFT	Fast Fourier Transform
FG	Focused Global, combine local and global information to improve the performance of models
FP	False Positive
FWHM	Full Width at Half Maximum
FQDN	Fully Qualified Domain Name
GOUS	Group Observation Units Set
GP	Gaussian process, powerful and flexible framework in statistics and machine learning
GPU	Graphics Processing Unit
GRU	Gated Recurrent Unit, type of RNN capable to learn from sequences of data
HFP	High spatial Frequency Patterns, for better interpretation and enhancement of images
HPC	High Performance Computing
IFT	Information Field Theory

INAF	Istituto nazionale di astrofisica
INFN	Istituto Nazionale di Fisica Nucleare
IoU	Intersection over Union, fundamental metric for evaluating the performance of object detection
IP	Internet Protocol
IR	Infrared
IRA	Istituto di Radioastronomia
JAO	Joint ALMA Observatory
JWST	James Webb Space Telescope
LSTM	Long short-term memory units, type of RNN designed to handle sequence prediction problems
MAP	Maximum-A-Posteriori
MEM	Maximum Entropy Method, reconstruct images from limited or noisy data
MFS	Multi-Frequency Synthesis deconvolution
ML	Machine Learning
MLS	Machine Learning Solution
MNRAS	Monthly Notices of the Royal Astronomical Society
MPA	Max Planck Institute for Astrophysics, Garching
MPI	Message Passing Interface
MS	measurement set
MTMFS	Multi-Term MFS deconvolution
NA	North America
NIFTy	Numerical Information Field Theory
NFMCP	New Frontiers in Mining Complex Patterns
NN	Neural Network
NRAO	National Radio Astronomical Observatory
ORP	Opticon Radionet Pilot
PBS	Portable Batch System
PDF	probability density function
PiP	Pip Installs Packages
PSF	Point Spread Function
QSO	Quasi Stellar Object
RADPS	Radio AStronomy Data Processing System
ResNet	Residual Neural Network, a popular CNN architecture
RESOLVE	Radio Extended SOurces Lognormal deconVolution Estimator [1, 18]
RMS	Root Mean Square
RPN	Region Proposal Networks, type of NN used in object detection tasks
RNN	recurrent neural network
SB	Scheduling Block
SED	Spectral Energy Distribution
Slurm	Simple Linux Utility for Resource Management
SNR	Signal to Noise Ratio
SSH	Secure Shell
STC	Scientific Technical Committee
SKA	Square Kilometer Array
SV	Science Verification data
SZ	Sunyaev-Zel'dovich effect
TAP	Table Access Protocol
TNG	Illustris The Next Generation project [21]
TP	True Positive
U-Net	“U” shape Network (specialized CNN for image segmentation)
UniFi	University of Florence, Italy
UniNa	University of Naples Federico II, Italy
VLA	The Karl G. Jansky Very Large Array
VLBI	Very Long Baseline Interferometry
VLTI	Very Large Telescope Interferometer
WSU	Wide Band Sensitivity Upgrade

1 Study evolution and achievements

The starting of this study coincided with the beginning of the pandemic and a significant slow down of the scientific world. Despite the outbreak, crucial advancements were made by the BRAIN study members to demonstrate that artificial intelligence is suited to tackle both detection of extended emissions and to provide speed-up procedures.

A well attended kick-off meeting occurred on December 11, 2020. The great affluence to this meeting demonstrated the importance to address current open questions in image analysis from a different perspective and technology. The goal of this study is to investigate advanced techniques, i.e. those algorithms capable to learn from the data and to overcome known imaging issues, such as detection of extended emissions and speed up convergence procedures. A core team was created with astro-informatics, astro-statistics, ALMA and `CASA` experts to address the goals of this study.

Following the requests during the initial meeting, seminars were provided by Kumar Golap, Tak Tsutsumi, Ben Bean (NRAO) and Dirk Petry (ESO) to address the `tCLEAN` task specifications, mosaicking issues and the ALMA data structure.

On December 15, 2022, the mid-term review meeting was held. The panel recommended continuing the study and advised engaging with the `CASA` and Pipeline teams. They also suggested further investigations, including: (1) addressing complex non-Gaussian source detection both spatially and spectrally; (2) applying the method to ALMA continuum and cubes; (3) enhancing faint source detection; and (4) exploring joint methods. All of these recommendations have been addressed, and the outcomes are detailed in this report.

A weekly work cadence has been established to maintain a sense of accomplishment and track progress. A dedicated group on Teams was created to document advancements, including meeting minutes, presentations, publications, and shared files.

The milestones of this study can be summarized as follows:

1. December 2020: **Kick-off** meeting
2. December 2020-February 2021: coordinated **lectures** on ALMA data structure and analysis
3. September 2021+: **ALMA SV data** applications
4. March 2022: Proposal submission to **C2PAP**, titled “Enabling Big Data Science in ALMA 2030 with machine learning” (Delli Veneri M., Guglielmetti F., Testi L.)
   - ∗ 3 computational nodes, 256 GB of RAM and 3 TB of storage
5. April 2022: application to **Leibniz Data Center** through C2PAP support
   - ∗ Cloud services: CPU : 1 computational node, 256 GB of RAM, 3 TB storage, GPU: 2 GPUs with at least 12 GB of memory, 480 hours for Training and Inference, between 720 and 2160 hours of fine tuning depending on the cluster availability
6. May-June 2022: Michele Delli Veneri is an ESO **visitor** for one month
   - ∗ Contributed talk at **SciOps2022**, titled “Data Cleaning, Detection and Characterization of Sources in ALMA Data through Deep Learning” (Delli Veneri M., Tychoniec Ł., Guglielmetti F., Villard E., Longo G.)
   - ∗ **ESO Azure MLS** accounts (Delli Veneri M., Tychoniec Ł., Guglielmetti F.) CL-PROD-001 MLS workspace. This is a Microsoft Data Science Solution for ML supported by ESO. MLS environment tested with our software, but on limited practice due to unsafe ownership on the software.
   - ∗ Contributed talk at **AI Forum**, titled “3D Source Detection and Characterization of Sources in ALMA Data through Deep Learning” (Delli Veneri M.)

- ∗ **ALMA data cubes simulation development and acceleration**. Available on GitHub with 100 ALMA data cubes (256x256x128) created in 20 seconds.
- ∗ `RESOLVE` and `DeepFocus` **applications** to simulated and real ALMA data.

7. July 2022:
   - ∗ **MaxEnt2022 conference**: Invited talk "Bayesian and Machine Learning Methods in the Big Data era for astronomical imaging" (Guglielmetti F., Arras P., Delli Veneri M., Enßlin T., Longo G., Tychoniec Ł., Villard E.) [22]
   - ∗ **MaxEnt2022 conference**: Contributed talk "Bayesian statistics approach to imaging of aperture synthesis data: RESOLVE meets ALMA " (Tychoniec Ł., Guglielmetti F., Arras P., Enßlin T., Villard E.) [23]
8. August 2022: article submission to **MNRAS** "3D Detection and Characterisation of ALMA Sources through Deep Learning" (Delli Veneri M., Tychoniec Ł., Guglielmetti F., Longo G., Villard E.) [2], published November 2022.
9. September 2022: **NFMCP (ECML-PKDD 2022)** workshop: Contributed talk "3D Detection of ALMA Sources through Deep Learning" (Delli Veneri M. et al.)
10. October 2022:
    - ∗ **ARC All–Hands Meeting**: Contributed talk "Update on the ESO internal ALMA development study: ALMA2030 Bayesian Reconstruction through Adaptive Image Notion" (Guglielmetti F.)
    - ∗ `RESOLVE` **Workshop 2022** at MPI for Radioastronomy. Contributed talk "Bayesian and Machine Learning Methods in the Big Data era for astronomical imaging" (Guglielmetti F. et al.)
    - ∗ **Comparison** of `DeepFocus` and `tCLEAN` on a set of 1,000 simulated ALMA data cubes.
    - ∗ First applications of `RESOLVE` to aggregate continuum and mosaicking in the search of SZ effect.
11. November 2022:
    - ∗ Contribution to **ORIGINS Annual Science Meeting** for GPU usage employing the C2PAP infrastructure.
    - ∗ Taking advantage of the *DRM F2F meeting* at NRAO (Charlottesville, VA): Presentation of the study to **Amanda Kepley (NRAO), Ryan Loomis (NRAO), Theodoros Nakos (JAO)** with a talk titled:"ESO internal ALMA development study: ALMA2030 Bayesian Reconstruction through Adaptive Image Notion" (Guglielmetti F.). Proposal from Ryan Loomis to make use of the new software in an hybrid system in `CASA`.
12. December 2022:
    - ∗ **EASC All-Hands meeting**: Contributed talk "Update on the ESO internal ALMA development study: ALMA2030 **B**ayesian **R**econstruction through **A**daptive **I**mage **N**otion" (Guglielmetti F.)
    - ∗ **Mid Term review process of this study**. The panel support the continuation of the study due to the high quality of the report and the work of the team. The panel recommends to 1) operate on sources with complex non-Gaussian structure, both spatially and spectrally; 2) ensure both methods work on continuum and cubes; 3) push the methods to fainter sources; 4) investigate how RESOLVE and Deep Learning techniques could be integrated into a joint method.
    - ∗ The team meets to discuss on the outcome of the mid-term review and starts with the requested advancements.
13. January 2023:

- Discussions commence on integrating `RESOLVE` and `DeepFocus` for enhanced functionality.
- **Co-Chair of the VLTI and ALMA Imaging Workshop**, highlighting the potential of `RESOLVE` and `DeepFocus` to address challenges in VLTI synthesis imaging.
  - Talks have been given by several contributors to this study with introductory talks on the topics of ML given by Prof. Giuseppe Longo and on Information Field Theory (IFT) given by PD Torsten Enßlin. Specific developments and applications have been discussed by Michele Delli Veneri, Łukasz Tychoniec, Philipp Frank, Jakob Roth and Jakob Knollmüller.
  - The Report on the ESO workshop: VLTI and ALMA Synthesis Imaging Workshop can be found at [24].

14. February 2023:

    - `ALMASim` has initiated new developments to tackle the challenges of simulating extended emissions in ALMA continuum and cubes. To improve accuracy, a more detailed characterization of noise is required, and an empirical noise study is being integrated to create synergies with `ALMASim`. Additionally, the algorithm is undergoing a significant reimplementation aimed at transitioning away from `CASA`.
    - The `RESOLVE` team designs new methods and applications to improve the speed up convergence procedures.

15. June 2023:

    - **Invited talk at the Statistical Challenges in Modern Astronomy VIII**: “A BRAIN study to tackle imaging with artificial intelligence in the ALMA2030 era” (Guglielmetti F.)

16. July 2023:

    - **MaxEnt2023 conference**: Poster contribution “A BRAIN Study to Tackle Image Analysis with Artificial Intelligence in the ALMA 2030 Era” (Guglielmetti F., Delli Veneri M., Baronchelli I., Blanco C., Dosi A., Enßlin T., Johnson V., Longo G., Roth J., Stoehr F., Tychoniec Ł., Villard E.) [25]
      - Guglielmetti F. part of the advisory committee for the MaxEnt2023

17. September 2023:

    - **STC ESAC online Meeting**: Invited talk “BRAIN” (Guglielmetti F.)

18. October 2023:

    - **ASTROINFORMATICS-2023 Conference**: (1) Invited talk **“A Brain Study to Tackle Imaging with Artificial Intelligence in the ALMA 2030 ERA”** Guglielmetti F. [26],(2) Contributed talk **“Deep Focus and ALMASim – A meta-learner for the resolution of inverse problems and its companion simulation package”** Delli Veneri M. [27]

19. November 2023:

    - **ORP Consortium Meeting**: Prof. Gerry Gilmore (Cambridge University/FORTH) has summarized the key outcomes of the BRAIN study.

20. December 2023:

    - **ALMA at 10 years: Past, Present, and Future**: Poster contribution with remote attendance “A BRAIN study to tackle imaging with artificial intelligence in the ALMA2030 era” [28]

- Drafting a DFG proposal with a group of collaborators to secure funding for enhancing `RESOLVE` applied to ALMA data and other ESO facilities.

21. January 2024:
    - ALMA Workshop 2023A, Ishigaki Island, JP: Invited talk “A BRAIN Study to tackle Image Analysis with Artificial Intelligence” (Guglielmetti F.)

22. February 2024:
    - **ESO Informal discussion**: “A BRAIN study to tackle imaging in the ALMA2030 era” (Guglielmetti F.)

23. April 2024:
    - **Submission Large Grant INAF Ricerca fondamentale “Enhancing Real-Time Image Analysis for SKA and ALMA Radio Interferometry”**: Our focus is on elevating `DeepFocus`, a Meta-Learner, for real-time image analysis. We propose to expand the capabilities of `ALMASim`, `DeepFocus`’ simulator, to tackle issues specific to new-generation facilities and provide real-data-based sky simulations suitable for training ML algorithms. `DeepFocus`’ improvements will allow us to mine archival data to search for serendipitous sources on known samples, such as pulsars. The software is built on an open-source framework and capable of being adapted to future facilities as AtLAST. Our collaboration brings together expertise from observatories such as ALMA and SKA, as well as AtLAST, encompassing proficiency in data mining and ML.
    - **Submission Mini Grant INAF Ricerca fondamentale “Towards an Empirical Noise Identificator and Generator for Multi-wavelength Astronomical Images”** (PI I. Baronchelli) to expand the capabilitites of `NOISEMPIRE`.:

24. May 2024:
    - **IRA Coffee talk, Bologna**: Invited talk “NOISEMPIRE, an empirical approach to ALMA noise simulation” (Baronchelli I.)
    - **IRA Coffee talk, Bologna**: Invited talk “Deep Focus: a Deep Learning Metalerner for ALMA imaging” (Delli Veneri M.)
    - **AI for Radioastronomy** Mini-Conference:
        - Invited talk “DeepFocus” (Delli Veneri M.)
        - Invited talk “A BRAIN study to tackle imaging in the ALMA 2030 era” (Guglielmetti F.)

25. June 2024:
    - **IRA Coffee talk, Bologna**: Invited talk “A BRAIN study to tackle imaging in the ALMA2030 era” (Guglielmetti F.)
    - **THE PROMISES AND CHALLENGES OF THE ALMA WIDEBAND SENSITIVITY UPGRADE** Conference: Invited talk “A BRAIN study to tackle imaging in the ALMA 2030 era” (Guglielmetti F.) [29]

26. September 2024:
    - Julian Rüstig (MPA/DZA) visiting ESO to contribute to the GOUS/ data combination imaging with `RESOLVE`

27. October 2024:
    - Final review of this study

28. November 2024:
    - **ARC All–Hands Meeting**: contributed talk “A BRAIN study to tackle imaging in the ALMA2030 era” (Guglielmetti F.)
    - **ORP consortium meeting**: Presentation on Joint Activity 3.4, focused on enhancing interoperability and improving user experience across optical and radio interferometry.
    - The final report of the BRAIN study, addressing all RIDs, has been submitted to the ESO Internal ALMA Development Study Coordinator.
29. December 2024:
    - **ADVANCED DATA PRODUCTS** Conference: Invited talk “A BRAIN study to tackle imaging in the ALMA 2030 era” (Guglielmetti F.)

2 The Core Team

Fabrizia Guglielmetti is ARC Scientist at ESO and Principal Investigator of this study. Her interest is on the development of advanced imaging techniques. She earned her PhD in astrophysics on image analysis employing Bayesian Probability Theory for a joint estimation of background and celestial signals (LMU, 2010).

Philipp Arras, formerly PostDoc at MPA, is expert on Bayesian statistics applied to imaging algorithms for radio telescopes. He developed theoretical models for inference algorithms for the unification of calibration and imaging, multi-spectral imaging, polarization imaging, data fusion with single dish data and Very Long Baseline Interferometry (VLBI). Expert on the `RESOLVE` technique.

Ivano Baronchelli (INAF/IRA) an expert in multiwavelength data reduction and analysis, with particular reference to the study of Galaxy formation and evolution in the context of large spectro-photometric surveys. He is currently specializing in ML techniques and he is the developer of `NOISEMPIRE`.

Matteo Bonato (INAF/IRA) an expert in ALMA data analysis and simulation generation, he extensively contributes to `ALMASim`.

Giuliana Cosentino is a Research Fellow at ESO, where she conducts her own scientific research using ALMA and other telescopes. She is expert in star and cloud formation and how these affect galaxy evolution. Part of her functional work is dedicated to this study.

Michele Delli Veneri (INFN) is an expert in astro-informatics and data science, with a master's in Astrophysics and a PhD from the Department of Electrical Engineering and Information Technology at the University of Naples Federico II. As the developer and designer of `DeepFocus` and `ALMASim`, he has gained extensive experience in imaging for both ALMA and the Square Kilometer Array (SKA).

Torsten Enßlin (MPA), Associate Professor at the Ludwig-Maximilians-University and IFT Group Lead. He is expert in image analysis for multi-wavelength observations, with special emphasis on radio synthesis observations.

Vishal Johnson (MPA), PhD student at LMU and expert in IFT, astro-statistics and `RESOLVE`.

Giuseppe Longo (UniNa), Professor in astrophysics, chair of the Data Science Initiative at the University Federico II: Data Science program, current president of the International AstroInformatics Association. His research interests are in development and application of novel ML methods to a variety of problems.

Jakob Roth (MPA) PhD student at LMU with experience in Bayesian inference algorithms and their application in astronomy, and especially radio interferometry. Jakob Roth contributed to the development of `NIFTy` and `RESOLVE`.

Alvi Rownok (UniNa) is a graduate student of UniNa, specializing in the development of astro-informatics techniques.

Julian Rüstig (MPA/DZA) is a PhD student at LMU with expertise in IFT, astro-statistics, and `RESOLVE`.

Luca Sannino (UniNa) is a graduate student at the University of Federico II and working on `ALMASim`.

Łukasz Tychoniec (Leiden Observatory) former Research Fellow at ESO. He is expert in star and planet formation, protostellar jets, protoplanetary disks, early planet formation, submillimeter interferometry employing mainly ALMA and the James Webb Space telescope.

Eric Villard, ESO/ALMA Advanced Data Product Scientist (ESO) involved with ALMA observatory in operations since 2010: Commissioning Scientist, System Astronomer, Deputy Head of Data Management Group, Head of the ALMA Array Performance Group.

Since the start of the study, several contributors moved to new positions. Philipp Arras (MPA) is a strategy consultant at Bain & Company. Łukasz Tychoniec (ESO) moved to Leiden Observatory, where he joined the Allegro node. Carmen Blanco worked within the IFT group (MPA) and applied `RESOLVE` to ALMA data for the detection of the SZ effect. Andrea Dosi completed his master from the University of Naples Federico II and supported the development of `ALMASim`. Luca Sannino (UniNa) graduated with a degree in Physics and has since started his Master's program.

This Study is also supported by Paola Andreani (ESO), Massimo Brescia (UniNa), Stefano Cavuoti (UniNa), Ed Fomalont (NRAO), Matteo Guardiani (MPA), Alessandro Marconi (UniFi), Federico Montesino Pouzols (ESO), Urvashi Rau (NRAO), Andy Strong (MPE), Udo von Toussaint (IPP), Crystal Brogan (NRAO), Sanjay Bhatnagar (NRAO), Allen Caldwell (MPP), Sandra Castro (ESO), John Hibbard (NRAO), Elizabeth Humphreys (ESO), Mark Lacy (NRAO), Dirk Muders (MPI for Radioastronomy), Martin Zwaan (ESO).

3 Software selection

This study represents a pioneering effort to apply artificial intelligence (AI) methodologies to ALMA data for imaging purposes, marking the commencement of a comprehensive feasibility assessment aimed at revolutionizing data processing in the field. By addressing two primary objectives (i.e., determining the applicability of AI methods for ALMA data processing and evaluating their effectiveness), we lay the groundwork for future advancements in astrophysical imaging.

Our approach is grounded in the complementary disciplines of **astrostatistics** [30, 31] and **astroinformatics** [32, 33]. **Astrostatistics** provides the rigorous statistical framework required to model, analyze, and interpret the complex and often noisy datasets produced by ALMA. This ensures that we can make reliable inferences about underlying astrophysical phenomena, even in the presence of incomplete or uncertain data. **Astroinformatics**, in turn, offers the computational tools and techniques required to handle large-scale astronomical datasets, such as ML, data mining, and advanced visualization methods. Together, these fields enable both efficient data management and scientifically robust results.

The synergy of these two fields is essential: **astroinformatics** enables us to handle data at the scale and complexity required for modern astronomy, while **astrostatistics** allows us to extract meaningful, scientifically valid insights from that data. Together, they provide the foundation upon which we can explore the potential of AI to enhance ALMA's imaging capabilities and drive future innovations in astrophysical research.

3.1 Selected Algorithms: `RESOLVE` and `DeepFocus`

We utilize two distinct software tools, `RESOLVE` [1] and `DeepFocus` [2], both selected based on extensive experience in their respective areas.

`RESOLVE`, grounded in astrostatistics, has been extensively applied in a variety of fields, including astrophysics, astro-particle physics, biology and medicine, under the guidance of Torsten Enßlin, a key figure in the development of IFT. This methodology has advanced numerous areas of research (e.g., [34–37]). `RESOLVE` has demonstrated its effectiveness in self-calibration and imaging with joint uncertainty quantification, particularly with VLA and VLBI data [15–17]. Ongoing developments include extending this algorithm to polarization imaging, covering Stokes parameters I, Q, U, and V. An overview on IFT applications can be found at this link.

`DeepFocus`, rooted in astroinformatics, is derived from the work of Giuseppe Longo, a key contributor and founder in the astroinformatics domain. The contributions coming from his team are encompassing a large variety of topics as data mining, ML for image analysis, development of computational tools for astronomical surveys, virtual observatories (see [38–41]), and leadership in education and outreach within the field (more information can be found clicking at this link). Some of their contributions outside astronomy can be found at this link. The `DeepFocus` algorithm is built on autoencoder architectures, which have been optimized for ALMA data through extensive testing and fine-tuning of various deep learning models, such as Convolutional Autoencoders (CAE), Variational Autoencoders (VAE), U-Nets, and Region Proposal Networks (RPN). The model selection process also incorporated cutting-edge methods like DeepGRU and ResNet, which were benchmarked against alternative architectures like Transformers and EfficientNets.

The integration of these advanced methodologies aims to significantly improve data analysis workflows and the overall user experience in ALMA data processing.

3.1.1 Advantages

`RESOLVE` and `DeepFocus` bring significant advantages in terms of the quality of data products, and they have the potential to largely improve the user data-processing experience both if used as a standalone tool and in combination with `CASA`.

`RESOLVE` has proven its value through successful applications with VLA data, making it a promising candidate for ALMA. It offers the additional benefit of joint uncertainty quantification and shows potential for advancing ALMA's data processing capabilities. Meanwhile, `DeepFocus` demonstrates its adaptability to ALMA's imaging challenges, particularly through the use of advanced architectures like 3D Deep Learning models to resolve imaging issues. When extended with Bayesian surrogate models in the taxonomy process, it enables uncertainty quantification. Furthermore, `DeepFocus` can be implemented for real-time image analysis, serving both operational needs and user support.

Together, these tools have the potential to revolutionize the processing of ALMA data, enhancing both accuracy and efficiency in imaging workflows.

3.1.2 Additional Techniques Considered

The original proposal also considered two additional algorithms: the Bayesian Mixture Model technique, an unsupervised learning method that provides robust background-source separation, and the Maximum Entropy Method (MEM) algorithm, which, while nonlinear and useful for correcting limited sampling, faces challenges with prior knowledge assumptions.

Firstly, the Bayesian Mixture Model technique [42] is an unsupervised learning method. The mixture model technique [42] is equipped with a defined model to explore the data and extract the required information from the data for a robust background-source separation. Based on Bayesian probability theory, the technique is capable to jointly estimate source signal and background, providing a multi-resolution analysis for the detection of faint sources. This technique for the use of ALMA needs further development and exploration, mainly driven by technology advancements, e.g. [43]. Please note that there is another technique worth of dedicating time and effort, that is described in [44]. The methodology described in [44] takes into account the representation of spectral signal in absorption/decrement and emission. The principles can be implemented within other techniques, as `RESOLVE`. The methodology [44] was designed as a natural extension of the Bayesian Mixture Model technique [42]. However, a dedicated human power is needed to achieve these developments. Secondly, the maximum entropy deconvolution (MEM) algorithm [45] is not further supported. MEM is a nonlinear deconvolution algorithm and useful to correct for the limited sampling of the u-v plane. However, it assumes a sky brightness of the sources as a prior knowledge. This algorithm showed the prior knowledge not adapting well to the data and hindering faint source detection. However, improvements in model description could make MEM competitive again, building on previous work [45–47].

The development of these methods will require further exploration and workforce resources when applied to ALMA, also in conjunction with data coming from other observatories (as JWST, Chandra, SKA, AtLast).

3.1.3 Future Potentials

As interdisciplinary approaches continue to advance, emerging methods offer new opportunities for improving imaging techniques. These novel approaches require rigorous comparison with established methods to ensure their effective integration into the broader scientific community. We propose conducting comprehensive quality assessments of several algorithms from the literature, comparing their precision, accuracy, data handling, and processing speed.

The following list represents current promising developments in enhanced image analysis and data mining for ALMA data. However, this list is by no means static (new techniques and algorithms are continually being published, and others may rise to prominence as the field evolves). As such, this list is a moving target, and future work will undoubtedly expand upon or refine these approaches.

Promising Image Analysis Methods Some of the following works represent promising developments for AI-enhanced image analysis in astronomy:

[48] (2024): Image dynamic 3D structures.

[49] (2024): High-performance computing (HPC) implementation for efficient least-squares estimation of sky intensities in radio astronomy.

[50] (2023): Deep learning models for deconvolution in interferometric radio data.

[51] (2023): Deep learning applied to ALMA solar data imaging.

[52] (2023): Regularized Maximum Likelihood approach to ALMA interferometric continuum imaging.

[53] (2023): A wideband, wide-field spectral imaging and deconvolution technique.

[54] (2022): Bayesian modeling package providing uncertainty quantification for image and source properties.

[55] (2020): Super-resolution imaging in ALMA data using sparse modeling techniques.

[47] (2018): High-performance GPU implementation of non-gridded Maximum Entropy Method (MEM) for imaging.

[56] (2018): Convolutional neural networks (CNNs) for detecting extended extragalactic radio sources.

[57] (2017): WSClean, a multiscale, multifrequency deconvolution algorithm.

[58] (2017): Matched filter techniques for weak signal detection.

[59] (2017): Clumpfinding algorithm for detecting weak signals in interferometric data.

[60] (2016): Clustering and pattern recognition for analyzing ALMA data cubes.

Promising Data Mining Techniques Several ML approaches hold potential for mining valuable insights from large astronomical datasets:

[61] (2023): CNNs for detecting specific continuum sources in the ALMA archive.

[62] (2023): ML surrogate models for simulating protoplanetary disk evolution, bridging theoretical predictions with observational data.

[63] (2023): Superresolution techniques for analyzing protoplanetary disk structures.

[64] (2022): ML methods to automate detection and analysis within protoplanetary disks.

[65] (2020): Data mining techniques for spectral line analysis.

3.1.4 Conclusion

The recent Nobel Prizes in Physics and Chemistry underscore the transformative impact of AI methodologies in scientific discovery. In particular, the laureates have been recognized for their pioneering work in either developing new AI techniques or leveraging AI to advance research in their respective fields. This global recognition aligns with our study's objective to revolutionize ALMA data processing by integrating cutting-edge AI approaches, further illustrating the profound influence AI has on modern science.

The integration of AI techniques into ALMA data processing represents a significant step forward in the field, setting the stage for continued advancements and exploration of new methodologies. As interdisciplinary approaches evolve, it is essential to conduct rigorous comparisons between established and emerging techniques to ensure effective utilization within the broader scientific community.

3.1.5 A Call for Acknowledgment in AI-Driven Discoveries

Our AI tools are designed to evolve and improve over time, ultimately reaching their full potential to empower the scientific community in making groundbreaking discoveries. As these tools are refined and increasingly adopted, we seek to ensure that proper recognition is given for their development. We envision these teams' contributions as an integral part of the discoveries enabled by these tools, and we believe that their involvement should be acknowledged in any resulting publications and breakthroughs. In essence, this is a call to protect the intellectual property of these AI innovations, ensuring that these teams are credited as the creators and developers of the methodologies that drive future advancements.

4 RESOLVE

The `RESOLVE` software is a versatile radio aperture synthesis tool that utilizes Bayesian principles [66] and is framed within the context of IFT [67,68]. It is supported by `NIFTy`, which offers prior models, variational inference algorithms, and capabilities for simulated data generation. `RESOLVE` has been refined for the use of ALMA data. Taking the u-v-w coordinates from the calibrated measurement set (MS) in the u-v plane, the Response operator R acquires the information of the dirty beam — R takes into account the frequency dependence of the beam. R is an integral ingredient for the reconstruction of the detectable celestial signal. The input data (or calibrated MS) d are modelled as a combination of celestial signal s corrupted by the dirty beam and by the noise n (systematic and random errors): $d = Re^s + n$. The process of synthesizing an image involves estimating the posterior probability density function (PDF) of potential true sky signal configurations arised by the Hamiltonian sampling: $P(s|d) = \frac{e^{-H(s,d)}}{Z(d)}$. More information on the foundations and description of the algorithm can be found at [1,67–70].
To efficiently access the information embedded in non-linear, high-dimensional probability distributions, the IFT team has developed and implemented new procedures within the `NIFTy` package. These include techniques such as Metric Gaussian Variational Inference (MGVI) [71] and Geometric Gaussian Variational Inference (GeoVI) [72], as well as the integration of Google's high-performance ML framework JAX [73], which supports automatic differentiation. While JAX enhances computational efficiency, the variational approaches introduced by MGVI and GeoVI offer fast and accurate results by approximating the posterior distributions' local Gaussianity. These techniques can be employed by `RESOLVE` to accelerate convergence, achieving a substantial speedup of up to 10x. However, further applications are necessary to fully quantify the impact of these advancements.

To summarize, `RESOLVE` aims to significantly increasing the quality of the reconstructed images which was successfully demonstrated in the publications listed above. These improvements so far came at the cost of an increased computational demands. The development of `fast-resolve` [13] has substantially reduced these complexities, achieving competitive reconstruction times without compromising image quality.

4.1 Proof of concepts

Working on the u-v plane, the long lasting weighting issue of ALMA visibilities was addressed and solved [23,25]. The successful application to HL-Tau data (SV data) with `RESOLVE` and the comparison with `CASA` provided the proof of concepts. In Fig. 2 (image on the right), the continuum image at 1.3 mm (233 GHz) applied to one spectral window with a width of 1.827 GHz (128 channels) is obtained with `RESOLVE`. The image on the left shows the original reconstructed IF image using the full band 6 and band 7 data [74].
The image, shown on the right of Fig. 2, is derived by the algorithm as representation of the posterior mean of the Hamiltonian sampling of the reconstructed celestial signal. A full set of posterior density function estimations of the celestial signal are provided and shown in Fig. 3. For each $P(s|d)$ a corresponding uncertainty map is estimated. At the converged sample, `RESOLVE`'s answer to the ill-posed problem of image reconstruction

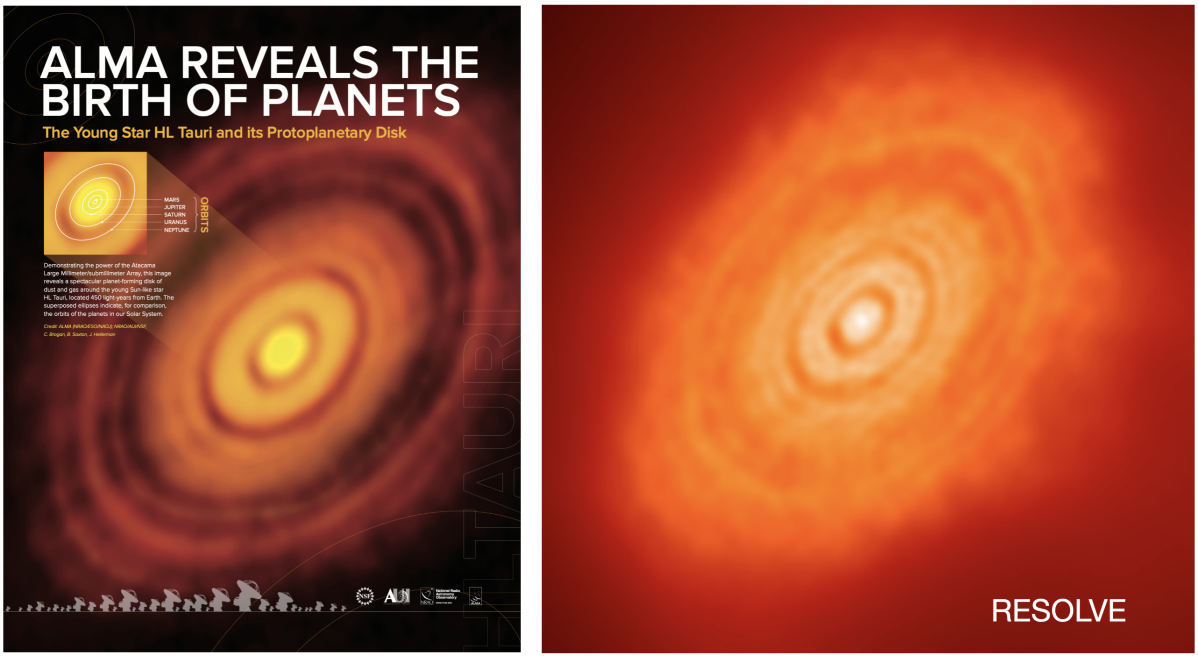

Demonstrating the power of the Atacama Large Millimeter/submillimeter Array, this image reveals a spectacular planet-forming disk of dust and gas around the young Sun-like star HL Tauri, located 450 light-years from Earth. The superposed ellipses indicate, for comparison, the orbits of the planets in our Solar System.

Credit: ALMA (NRAO/ESO/NAOJ); NRAO/AUI/NSF, C. Brogan, B. Saxton, J. Hellerman

Figure 2: On the left, HL-Tau press release combining bands 6 and 7 [74]. On the right, HL-Tau imaging with `RESOLVE` *on band 6 and one spectral window (one quarter of the available data for band 6).*

provides a posterior mean of the sampled estimated signals and a mean uncertainty map: See Fig. 4, on the upper/lower left are shown the posterior mean of the signal detection and the corresponding uncertainty map. Still in Fig. 4, the power spectrum $P_s(\vec{k})$ of the process that generated the signal s as a function of spatial frequency k and its uncertainty estimation are shown. A representation of the distribution of data weights versus baselines is also reported (right).

In the future, we plan to apply `RESOLVE` to HL Tau data by combining observations from Bands 6 and 7, comparing the results with those obtained using `CASA`, as shown in Fig. 2 (left image). Additionally, several SV datasets for HL Tau are available in Bands 3, 4, and 9, offering extremely high angular resolution. These datasets present a valuable opportunity for a comprehensive quality assessment and comparison of imaging results generated by `RESOLVE` and `CASA`. To fully realize this potential, further development of `RESOLVE` within the ALMA framework is essential, particularly to enhance support for self-calibration processes. Building on the foundational work of [15], [16], and [17], this advancement will require dedicated resources to refine the algorithm and rigorously test it on the available data. Securing funding for this effort will enable the integration of cutting-edge methods, ensuring significant improvements in imaging quality and calibration efficiency. For more details on SV data, see the ALMA Science Portal.

4.1.1 The power spectrum

The signal's power spectrum $P_s(\vec{k})$ describes how the signal's variance is distributed over the different frequencies of the signal [67]. Low- and high-frequency modes correspond to large- and small-scale features of the signal, respectively. An example of power spectrum reconstructed from the `RESOLVE` procedure for image reconstruction is shown in Fig. 4.

The power spectrum contains interesting information about the statistical properties of the physical signal.

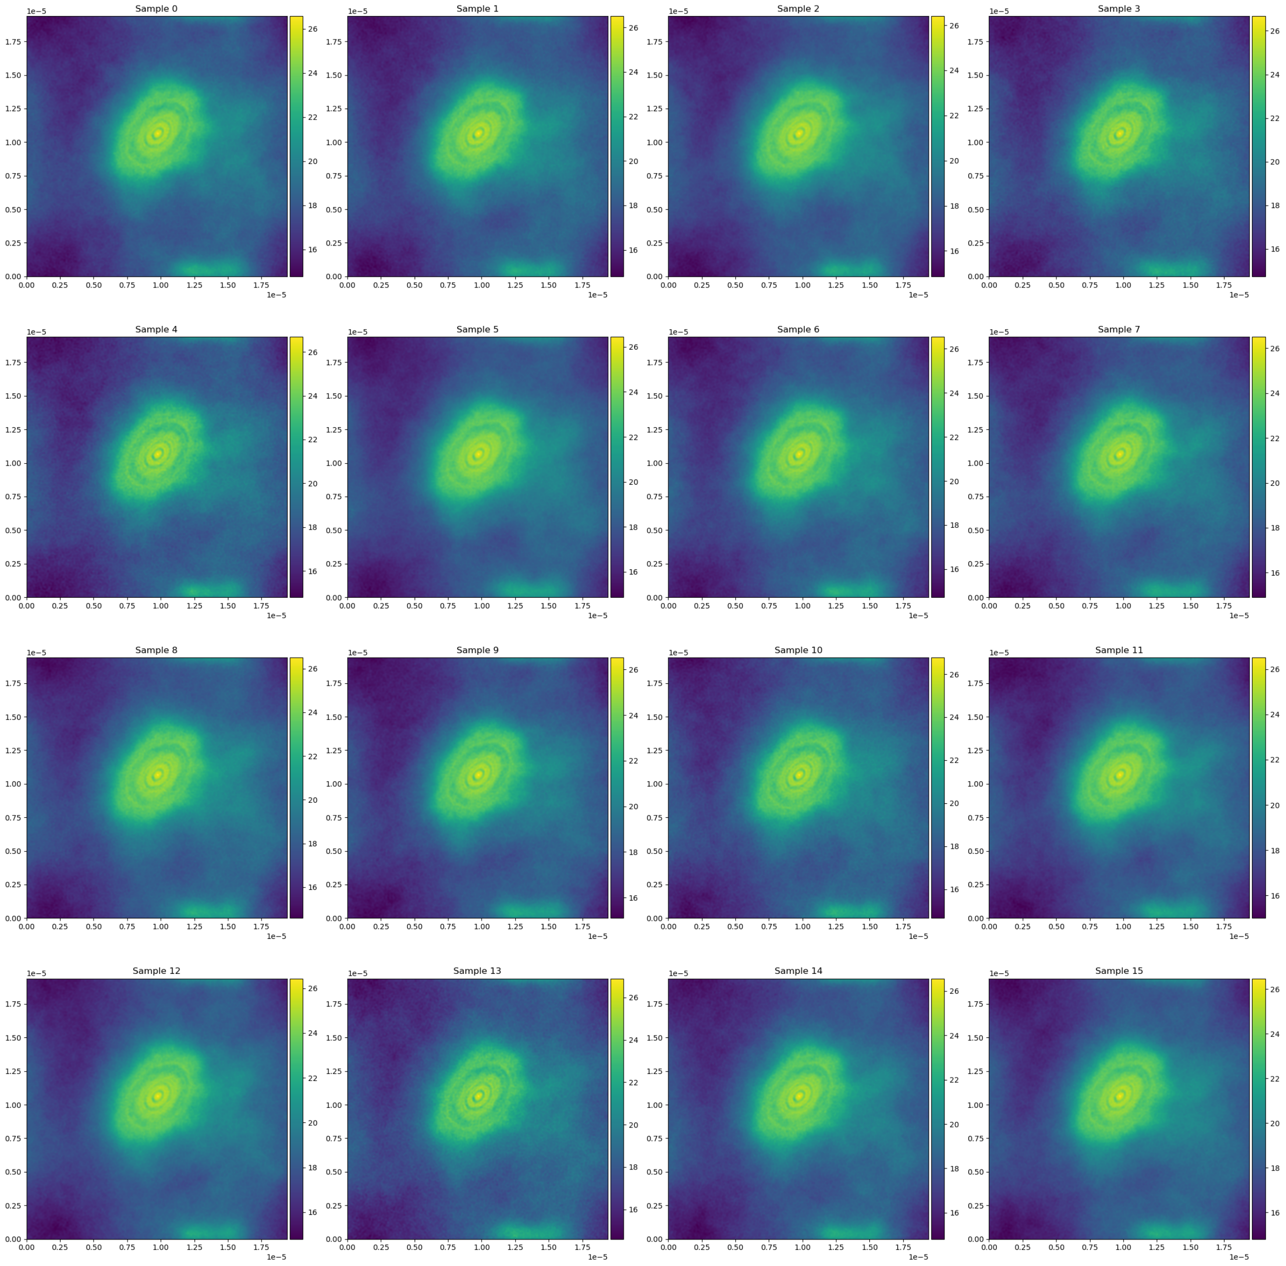

Figure 3: RESOLVE: Set of Hamiltonian samples displaying the estimated posterior PDF of the inferred celestial signal for the HL-Tau data shown in Fig. 2, right. Each sample is similar with negligible deviations in reconstructed intensity and shape, as they correspond to the final iteration of the converged optimization procedure.

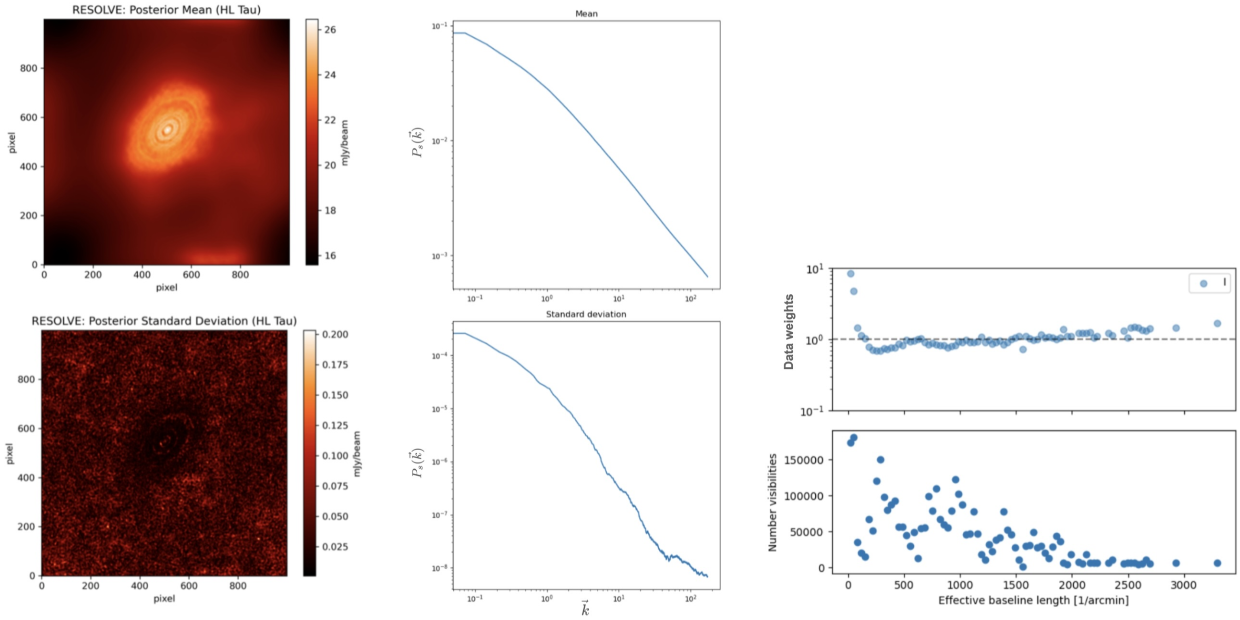

Figure 4: `RESOLVE`: Left, the mean posterior PDF (upper) and its uncertainty (lower) of the estimated celestial signal for the HL-Tau data shown in Fig. 2, right. Center, power spectrum estimation $P_s(\vec{k})$ corresponding to the estimated posterior mean (upper) and its uncertainty (lower). Right, representation of `RESOLVE`'s distribution of data weights versus baselines (upper) and number of visibility versus effective baseline length (lower) at the convergence iteration.

If the power spectrum falls with higher frequencies, it corresponds to higher variations on large-scale features compared to small-scale features. This is a typical behavior for many physical signals. When analyzed with sufficient resolution, such signals are smooth functions. The slope of the power spectrum determines the degree of smoothness. Steep spectra correspond to smooth signals, while flatter spectra correspond to signals with more small-scale variations. Peaks and bumps in the power spectrum indicate oscillations with the respective frequency. Flat power spectra correspond to white noise.

4.2 Simulations

The major advantage in using simulated datasets for quality assessment of the imaging procedures is the selection of the input parameters, allowing to control the shape and brightness of the sources, as well as configuration and behaviour of the telescope array. The downside of this approach is the difficulty to realistically model the noise acquired during an observation and the simplistic assumptions about the sky brightness. Nonetheless, this test allows to understand how `RESOLVE` works and to compare source detection and flux estimates with the `tCLEAN` task [12] in `CASA` [11, 20]. In the following, two applications are shown that are tailored at testing the algorithm capabilities to overcome source confusion and weak signal detection in single pointing and continuum images. The `RESOLVE` algorithm is shown to be successfully applied to ALMA simulated data.

The simulation in Fig. 5 aims at challenging the algorithm to detect weak and extended as well as bright point–like sources spread over an ALMA synthetic observation. The simulated data set is processed with `RESOLVE` and the `tCLEAN` imaging task in `CASA`. Taking the simulated data set as benchmark, `RESOLVE` provides both weak and extended source detection and more realistic source characteristics with respect to `tCLEAN`. `RESOLVE` powerfully provides detection of diffuse emission, weak signal, point–sources embedded in diffuse emissions.

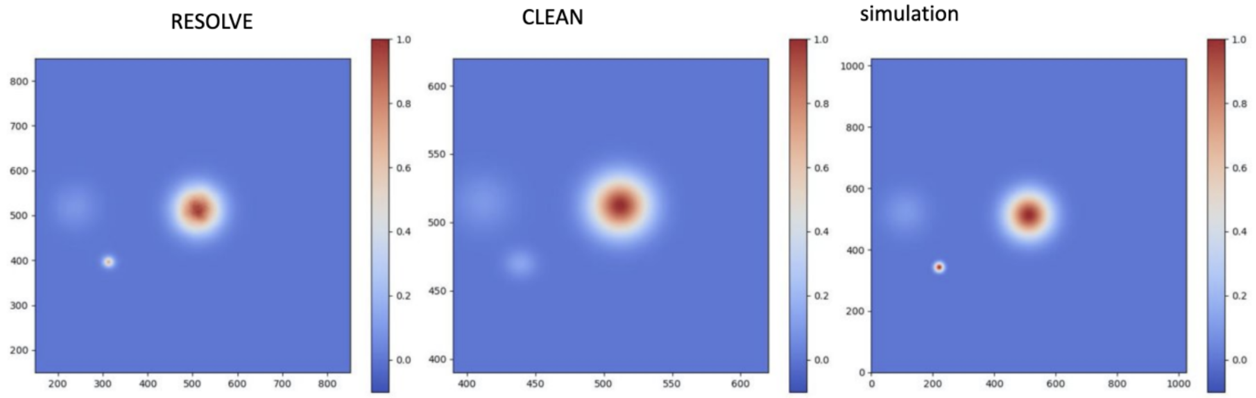

Figure 5: *`RESOLVE` (left) and `tCLEAN` (center) application to an ALMA continuum simulated data set (right). The colorbar shows a normalized flux intensity to allow for a visual comparison on recovered source characteristics.*

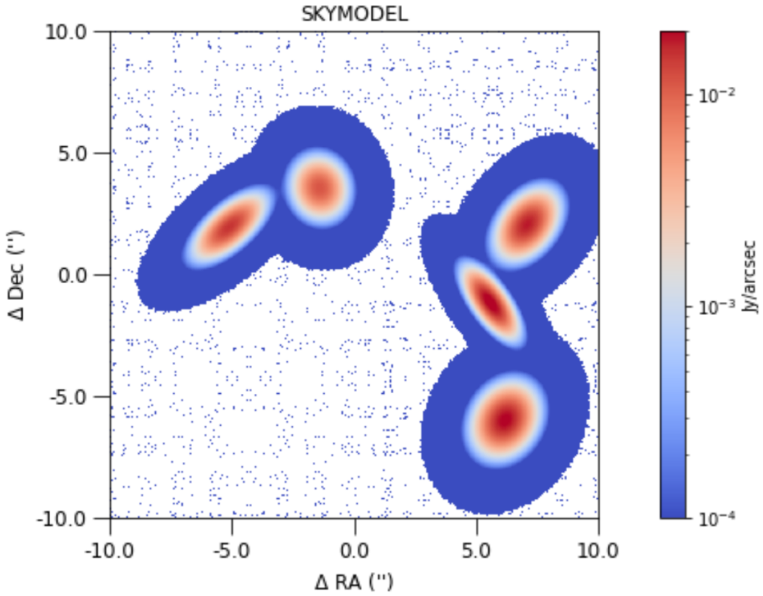

Figure 6: *Sky model which serves as an input for creating simulated ALMA MS: Single pointing and continuum image.*

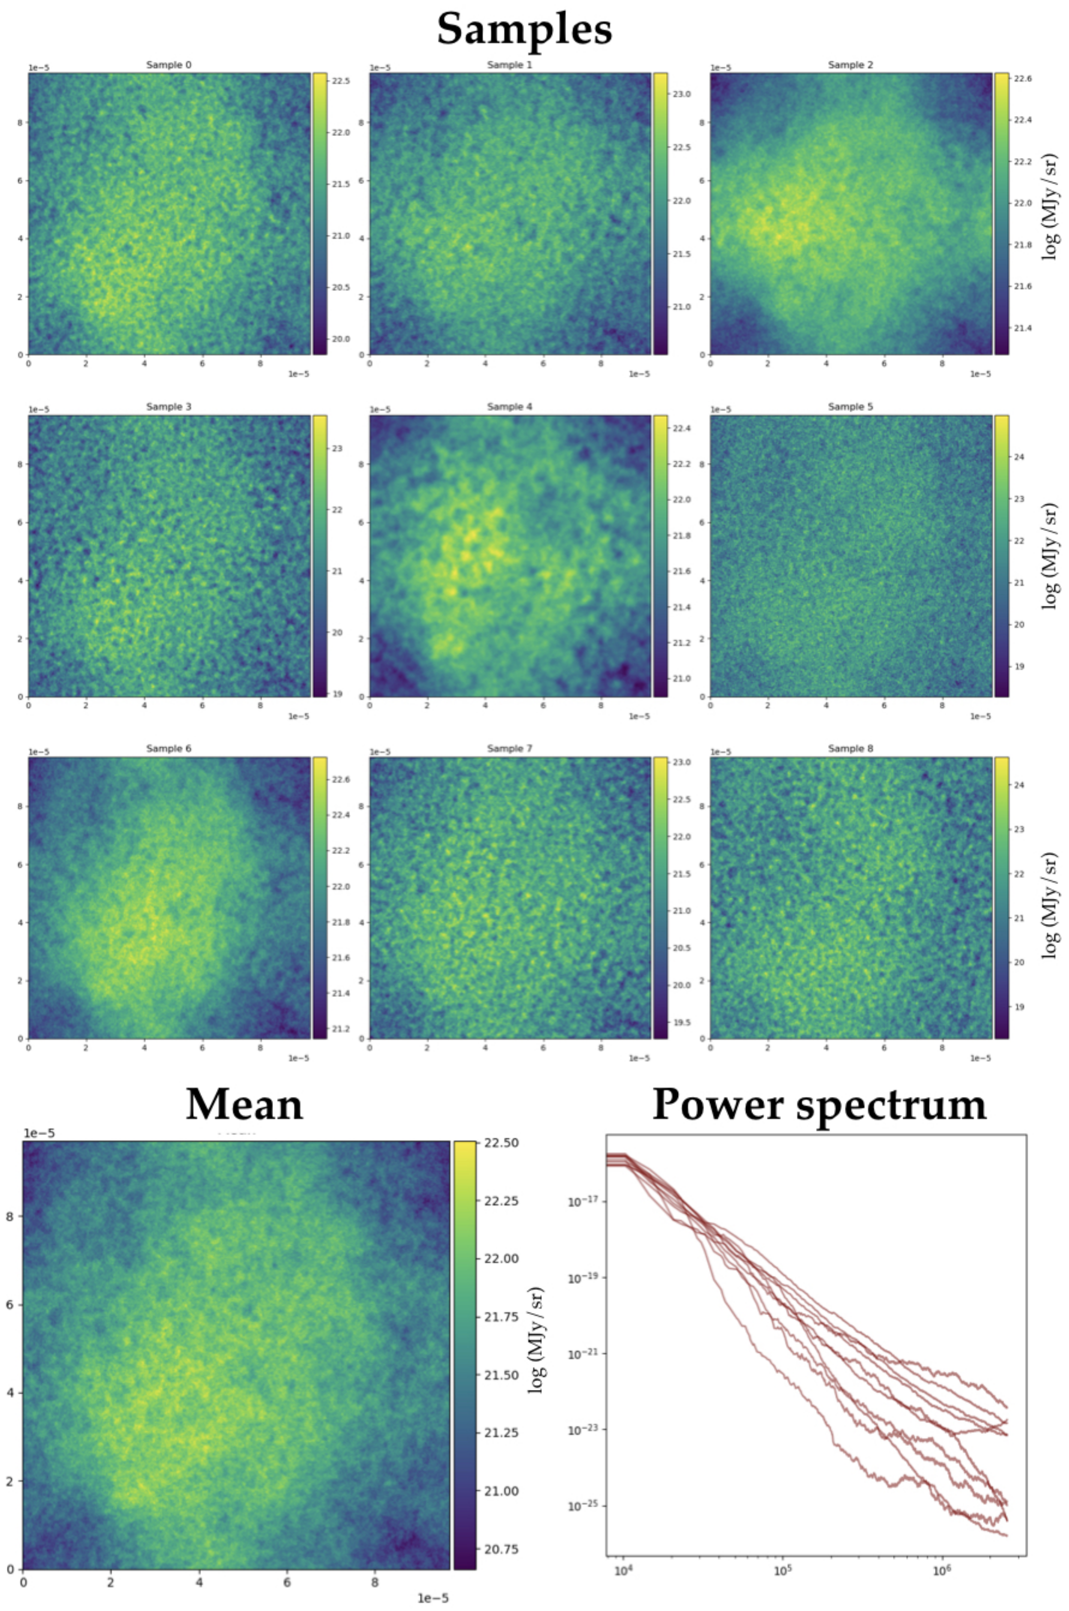

Figure 7: ***Iteration 0****: First set of Hamiltonian sampling from the starting iteration of ALMA single pointing and continuum simulated data set (Fig. 6). The initial uninformative guess of* `RESOLVE` *is shown.*

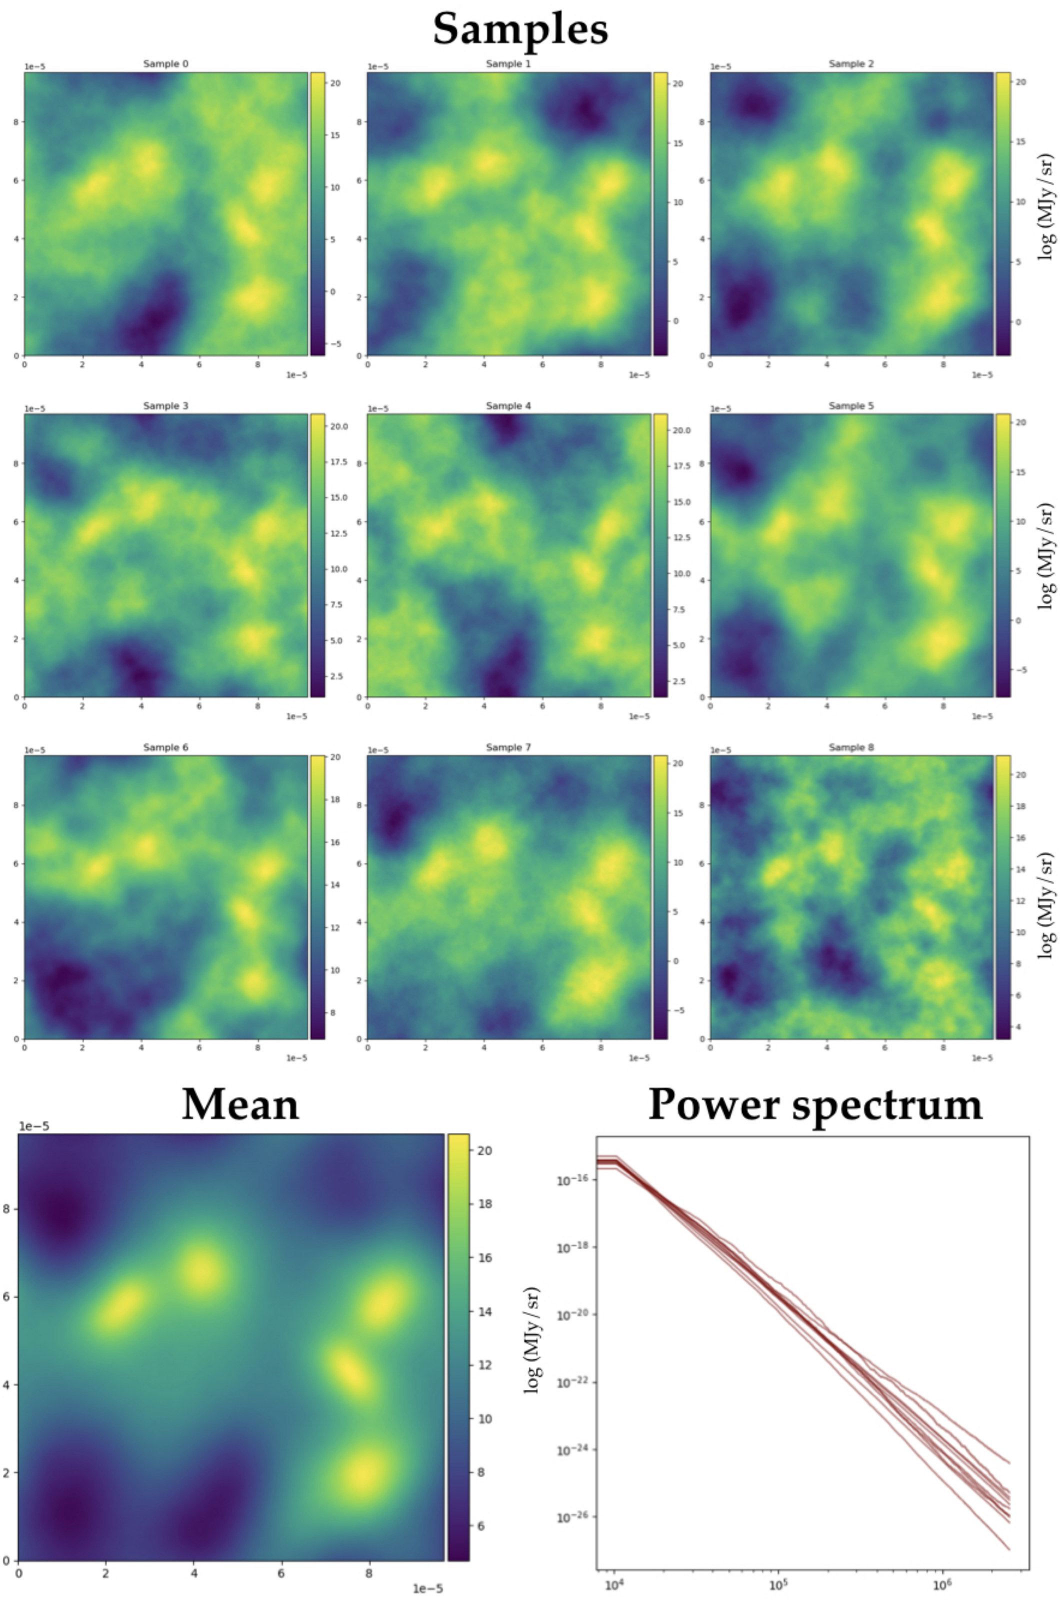

Figure 8: ***Iteration 5****: After 5 iterations, the set of Hamiltonian sampling of the ALMA single pointing and continuum simulated data set (Fig. 6) shows that the algorithm learns about the signal s.*

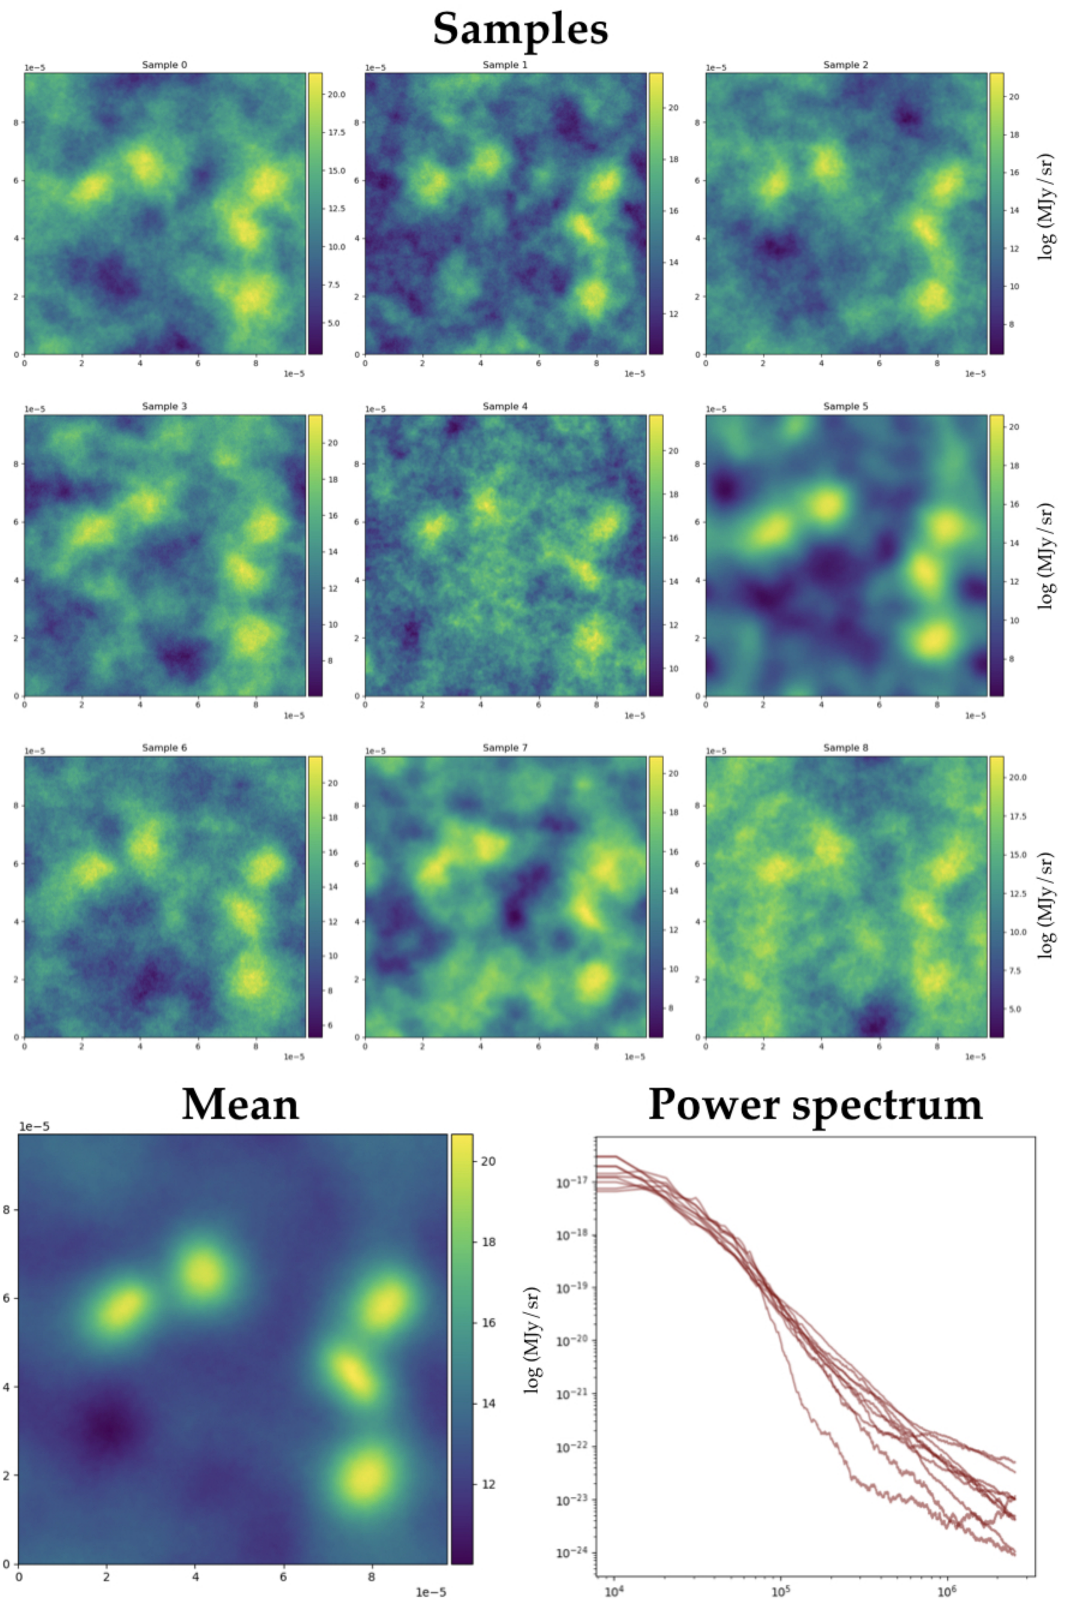

Figure 9: `RESOLVE` *learns the signal s (Fig. 6) as shown at (almost) convergence stage (**iteration twenty-four**). The mean of the posterior PDFs of the detected signal s given the data d is shown in the lower left image. The image on the lower right shows the reconstructed power spectrum samples, starting to deviate from the initial uninformative guess.*

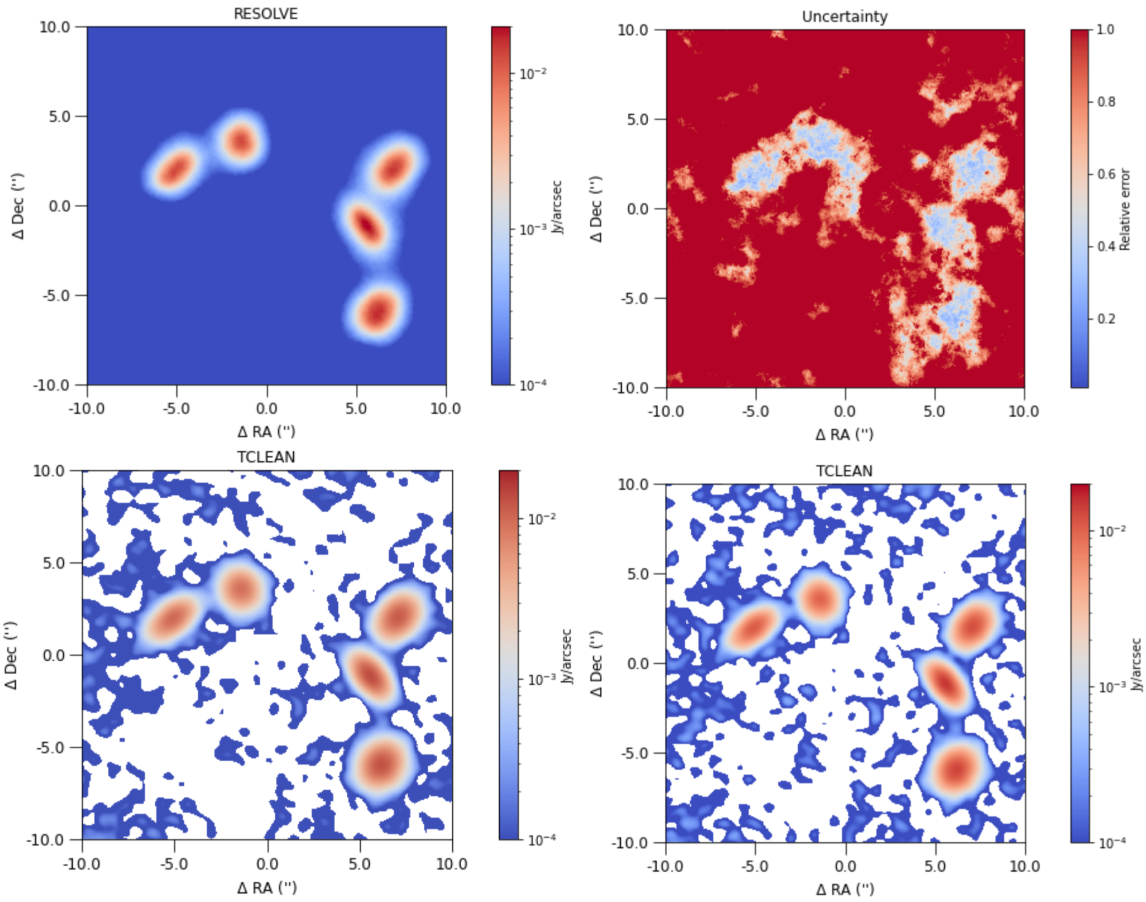

Figure 10: Maps of the brightness distribution in the sky plane of the simulated data (Fig. 6). Upper row: `RESOLVE`'s reconstructed image (left) and its uncertainty map (right). Lower row: `tCLEAN` solution using natural weighting (left) and robust = 0.5 following the Briggs weighting scheme (right).

Parameter	Mean
Offset	26
Zero mode	1±0.1
Fluctuations	5±1
Power spectrum slope	-2±0.5
Flexibility	1.2±0.4
Asperity	0.2±0.2

Table 1: Example of input parameters for the `RESOLVE` *application to simulated data and its estimated values.*

Fig. 6 shows the sky model generated with `simalma` task in `CASA`. A simple 2D array consisting of five Gaussian components of different brightness, size and position angle are generated. The sky model is generated, imposing physical properties to the sky, such as spherical coordinates, pixel dimensions, field-of-view and brightness in physical units. Afterwards, the task is simulating observations of the given sky model with the ALMA observatory. For this specific case we simulate observations with ALMA configuration C-3 at 230 GHz (ALMA band 6), which results in an effective resolution of 0.7 arcsec, due to the longest available baseline being ~ 500 m. The `simalma` task returns a calibrated MS, that consists of complex visibilities. Those visibilities are a Fourier transform of the sky brightness, therefore inverse Fourier transformation provides a dirty image of the observed sky.

In Fig. 10, lower row, the reconstructed images with `tCLEAN` are shown. The task `tCLEAN` is run without any constraint on where to look for point sources in the image (i.e., without any masking as input parameter), for 20×10^3 iterations, or until threshold of 0.3 mJy/beam is reached. The pixel size of the reconstructed image is set to 0.1 arcsec and the image size is set to 512×512 pixels. The noise threshold level of 0.3 mJy/beam is selected so that the `tCLEAN` algorithm does not attempt to find sources from the residual image consisting purely of noise. The RMS of the dirty image is 0.1 mJy/ beam and the SNR is 3σ. In the default settings, the weighting of the baselines is set to $natural$ (image on the left), which means it associates the baseline with weight proportional to the sampling density (i.e., the most covered baselines have the highest weight). Since at larger scales the sampling is much denser, this puts more weight on lower resolution, which results in achieving lower resolution than the sky model image. Natural weighting will give the largest beam and the best surface brightness sensitivity. Therefore we also attempt imaging with Briggs weighting [75] with robust parameter 0.5 (image on the right), which moves the balance of weighting toward longer baselines increasing the resolution but decreasing signal-to-noise ratio.

For `RESOLVE` image of the simulated data, 30 iterations are needed to reach convergence (see [13] for the latest information about clock times). In Figs. 7, 8, 9, the first, sixth, twenty-fourth iterations of `RESOLVE` are shown. The $P(s|d)$ samples are displayed, including the posterior mean $< P(s|d) >$ and the power spectrum samples $P_s(\vec{k})$. Given an initial state of no information, the algorithm searches for the optimal sky configuration from the given data (visibilities and dirty beam). In Fig. 10 (upper row), the converged posterior mean PDF of the detected simulated signal $< P(s|d) >$ and its uncertainty map are shown. The uncertainty map displays the relative error as a measurement of precision, to determine the magnitude of the absolute error in terms of the actual size of the measurement process. No assumption was made on the presence of point sources in the data. Input parameters of the `RESOLVE` run are summarized in Table 1: See [76] for a detailed explanation of the parameters.

`RESOLVE`'s reconstructed image exhibit a smooth background in agreement with the simulation (Fig.6). Bayesian probability theory allows us to estimate the uncertainty of the hypothesis of detecting a source signal. What is not a celestial source and potentially noise is captured by the description of uncertainty on the full image. Therefore, the uncertainty map does not only address detected sources. Structures in the image away from the detected sources provide evidence that the uncertainty quantification can not be a smooth function. The `tCLEAN` reconstructed images are characterized by negative values in several areas. Negative values arise

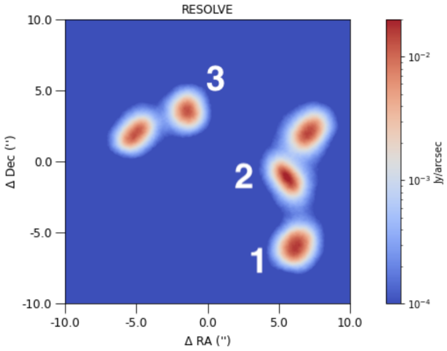

Comp	model	tCLEAN natural	tCLEAN Briggs	RESOLVE
1	34.11	24.67	24.96	26.68±5.76
2	23.31	21.68	22.19	23.59±6.72
3	16.69	16.31	16.57	17.82±3.86

Figure 11: Example of reconstructed fluxes (sources $1-3$) from the application of RESOLVE *and* `tCLEAN` *to Fig. 6*

Table 2: Peak flux (in Jy $arcsec^{-2}$) of different components on the simulated image.

by the recurrent application of subtraction during the major-minor cycles and residual estimations. Images obtained by `tCLEAN` and other `CLEAN` based algorithm often contain negative values in regions with very low or zero surface brightness. Physically this is not possible as there is no negative flux.[7] Moreover, a comparison of the reconstructed fluxes $[Jy \cdot arcsec^{-2}]$ with the two algorithms is shown in Table 2 where the corresponding components numbered 1-2-3 are indicated in Fig. 11. The model column provides the simulated sky model's flux values. Both algorithms provide flux reconstruction close to the ideal values. `tCLEAN` with $natural$ weigthing provides the worst estimates, being those simulated sources mainly point-like sources. `RESOLVE` gives reconstructed flux values (and their uncertainties) 7% closer to the simulated values with respect to `tCLEAN`.

`RESOLVE` within IFT enables the estimation of parameters along with their associated uncertainty distributions. While classical approaches do not inherently provide full uncertainty distributions, these can be derived as a second step on the image products. This additional analysis has not yet been considered in the current study but will be prioritized in future applications. For scenarios where the sky model predominantly comprises extended Gaussian structures, multi-scale `CLEAN` may provide a more computationally efficient and specialized alternative. Such cases will be carefully considered in future imaging verification studies.

4.3 DSHARP Large Program

The application of `RESOLVE` to real data is performed on a well-studied data set, whose survey was designed to optimize the spatial resolution and contrast sensitivity to continuum emission substructures [7]. This is the ALMA Large Project 2016.1.00484.L, known as Disk Substructures at High Angular Resolution Project (DSHARP). The sample is shown in Fig. 12. The data were taken at long baselines allowing us to test `RESOLVE` on high resolution data while detecting extended emission. In fact, a secondary goal of this project was to identify corresponding gas structures and infer other relevant bulk disk properties (e.g. geometry). Expertise on this science topic is brought into the study by Ł. Tychoniec (ESO).

The protoplanetary disk Sz114 is chosen within the sample, because of its relatively smooth structure [23]. Sz114 is shown as first image on the left, second row, in Fig. 12.

Both `tCLEAN` and `RESOLVE` are applied to a single spectral window to directly compare the two techniques. From the publicly available calibrated MS, we extracted spectral window 9 with split task in `CASA` and binned it into a single channel.

---

[7] Please note that physical values of the sky distribution can not be a negative number. Bayesian analysis allows one to assign, e.g., to the background component the signal not described by the celestial sources. It allows to correct for effects as the one due to the missing short spacing. If the prediction is consistent with the data, then an increment in the uncertainty map should also increase e.g. at correspondingly to the missing short spacing.

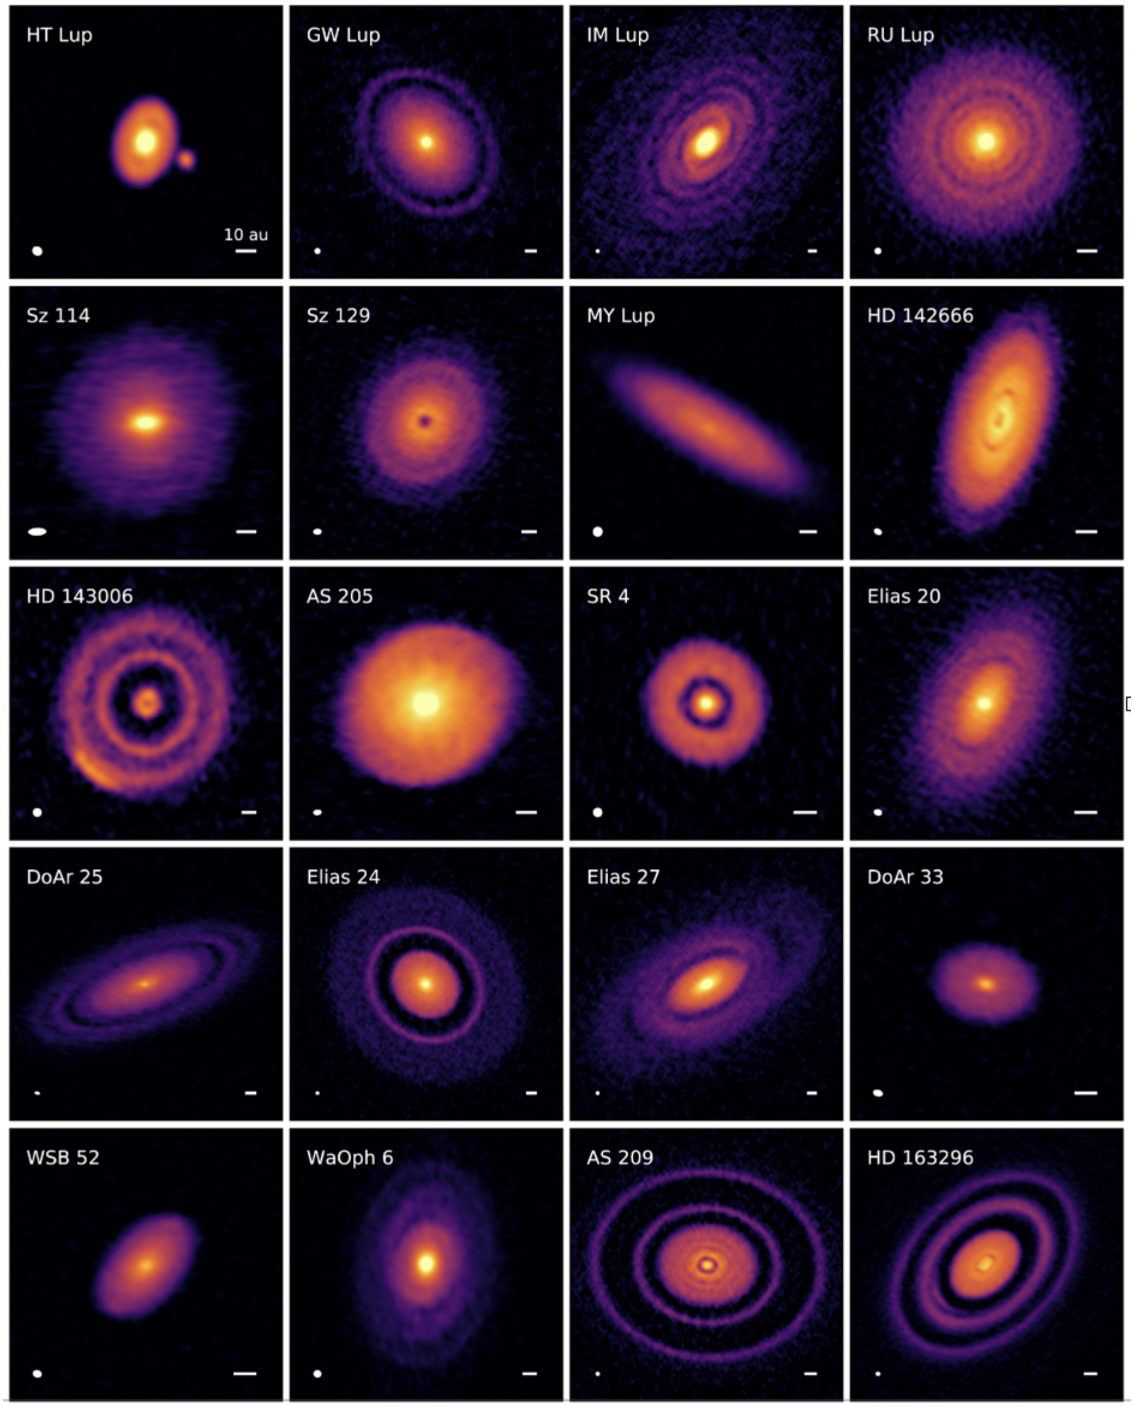

Figure 12: This image is taken from [7]: Gallery of 240 GHz (1.25 mm) continuum emission images for the disks in the DSHARP sample. Beam sizes and 10 au scalebars are shown in the lower left and right corners of each panel, respectively.

Parameter	Mean
Offset	20
Zero mode	1±0.2
Fluctuations	3±1
Power spectrum slope	-4±2
Flexibility	4±0.8
Asperity	2±0.8

Table 3: *Example of input parameters for the `RESOLVE` application to DSHARP data's Sz114 and its estimated values.*

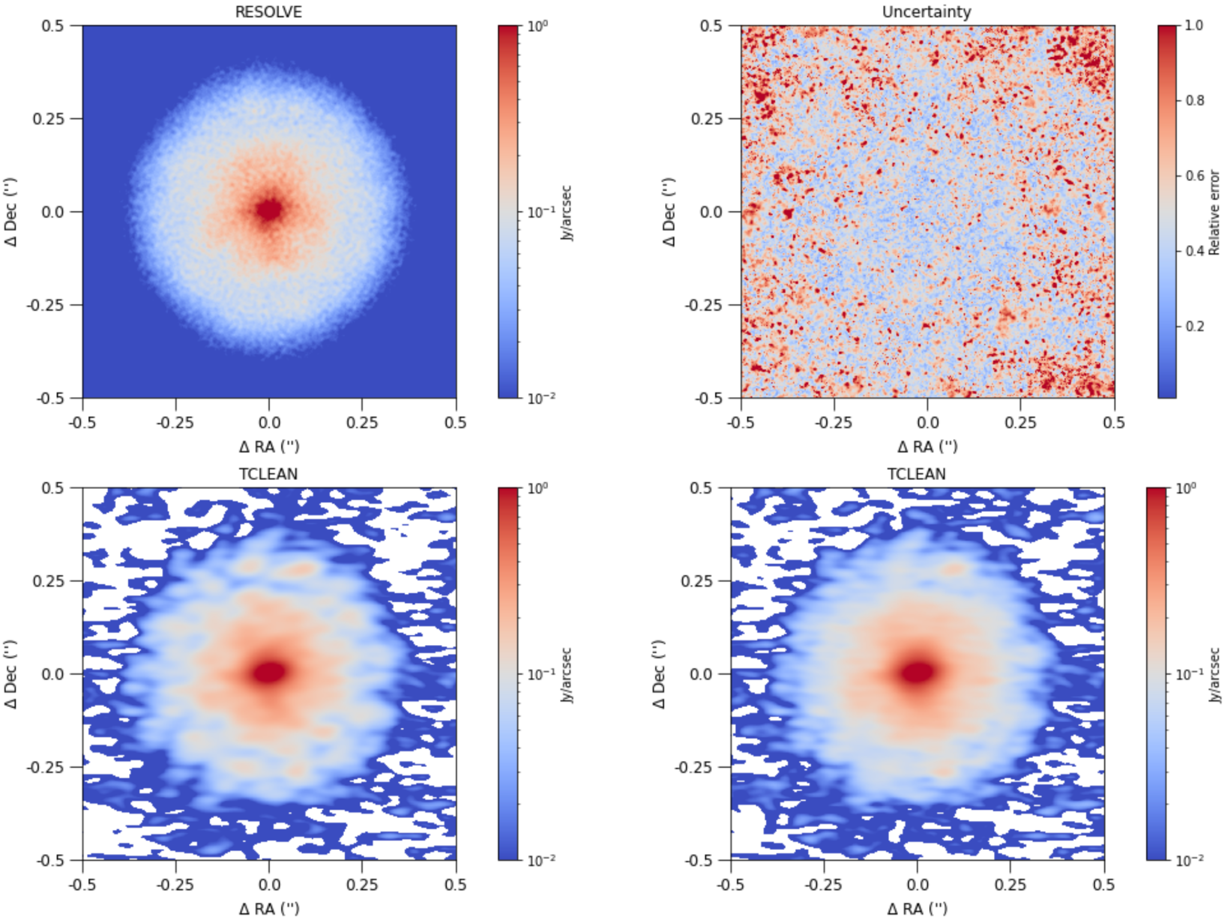

Figure 13: *Image reconstruction with `RESOLVE` and `tCLEAN` of ALMA observation of Sz114 (Fig. 12). Upper row: reconstructed image by `RESOLVE` (left) and the relative error map (right). Lower row: `tCLEAN` reconstructed images with standard cleaning algorithm (right) and with multi-scale algorithm (left).*

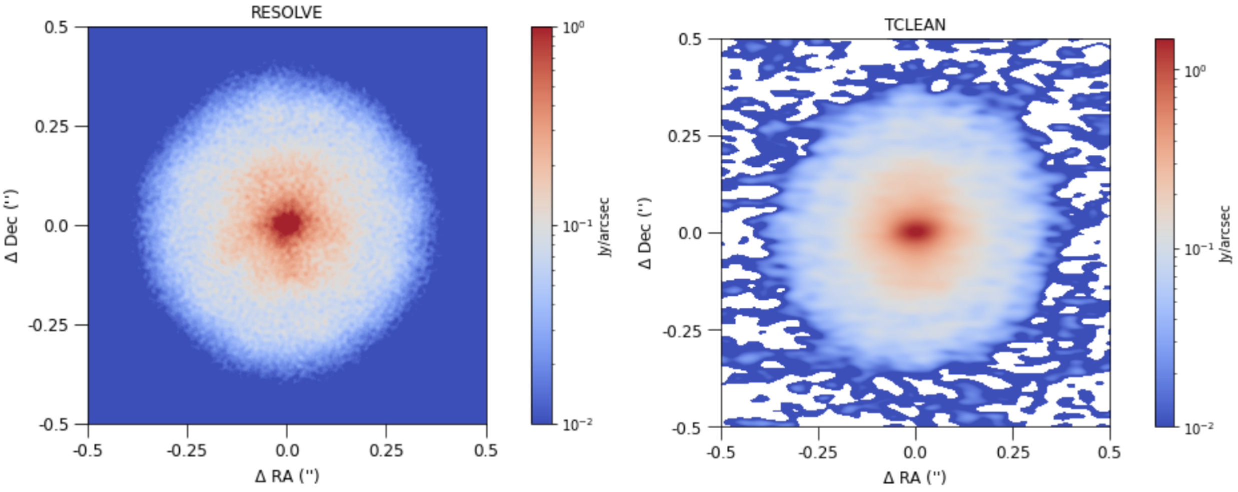

Figure 14: On the left, the reconstructed image by `RESOLVE` of Sz114 as shown in Fig. 13. On the right, the reconstructed image of Sz114 with `tCLEAN` as delivered by the Large Program team [7] (In table 4 indicated with "best tclean").

Radius (arcsec)	Best Fit	tclean		RESOLVE
		Hogbom	Multiscale	
Peak Flux (Jy·arcsec^{-2})	1.57	1.55	1.49	5.8 ± 3.82
0.06	8.63	8.85	8.95	9.48 ± 1.07
0.15	22.79	22.82	22.53	23.21 ± 0.68
0.35	47.45	47.07	47.24	47.24 ± 2.79

Table 4: Peak (in Jy·arcsec^{-2}) and integrated flux at different radii of the Sz114 disk.

The estimated input parameters of the `RESOLVE` application to Sz114 is shown in Table 3. Please note the power spectrum slope with respect to the result in Table 1. The achieved steeper power spectrum slope on the real application is reasonable since less structures are present at small scales with respect to the simulated case. This result is in agreement with the reconstructed image of Sz114 by the `RESOLVE` algorithm (Fig. 13), upper left, and the uncertainty map, upper right. The uncertainty map shows that no aggregate structures are detected and that the quantified uncertainty increases at increasing signal detection (absolute uncertainty).

Standard and multiscale `tCLEAN` imaging are used because the observed disk presents large variety of spatial scales. Images are characterized by 0.005 arcsec (pixel size), 1024×1024 pixels and 20×10^3 iterations or until the noise threshold of 0.05 mJy is reached. In case of multiscale clean, three scales (0, 7, 28 pixels) are specified to find three types of sources in the data: point source (scale=0), extended (Gaussian) components with FWHM of 7 and 28 pixels. The resulting `tCLEAN` images are displayed in Fig. 13, lower row.

In Fig.14, a visual comparison between the reconstructed images of Sz114 from `RESOLVE` and from `tCLEAN`. This `tCLEAN` image is produced employing the aggregate continuum as delivered by the Large Program team [7]. We define this `tCLEAN` solution on the aggregate continuum as the best `tCLEAN` image. `RESOLVE` is capable to achieve improved emission detection. Table 4 provides a comparison of integrated flux estimates measured over different areas of the disk (radius). The three `tCLEAN` solutions are very similar. The fluxes provided by `RESOLVE` are close to the benchmark values. The only deviation found is at the peak. `tCLEAN` approximate the dirty beam to a Gaussian, while `RESOLVE` encodes the dirty beam from the uv-plane into the Response operator. The distribution at the peak collects more flux than a Gaussian. In other words, it is expected that convolving a `RESOLVE` image with a Gaussian kernel of the size of that used by `tCLEAN` should produce images with the same peak flux — the discrepancy is hypothesized to be due to higher resolution reconstruction by `RESOLVE`. Therefore, it is reasonable that at the peak the measured flux is larger than the value provided by `tCLEAN`.

`RESOLVE` is providing an high-fidelity image of the protoplanetary disk Sz114 and robust estimation on the fluxes and their uncertainties.

4.3.1 Application to Elias 27: Continuum and Spectral Cube Imaging

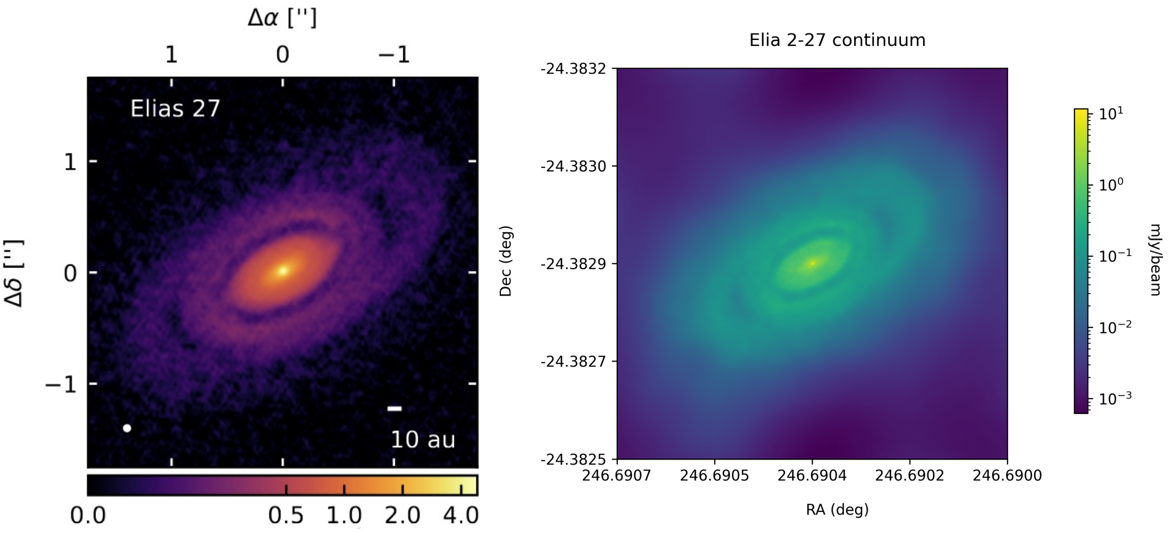

Figure 15: Elias 27 from the DSHARP ALMA project at 240 GHz (1.25 mm) continuum. On the left, the fiducial image as given by the DSHARP team [77]. On the right, `RESOLVE` mean sky map of Elias 27.

The DSHARP project sought to observe a large number of protoplanetary disks at high resolution (35 mas) to uncover small-scale substructures within the disk material and investigate their connection to the planet formation process. To assess the likelihood that detected structures are genuine, it is crucial to compare the results using multiple robust algorithms.

In Figs. 15 and 16, the application of `RESOLVE` to a protoplanetary disk sample, Elias 27, from the ALMA DSHARP project [7,77] is shown. In Fig. 15, the image on the left is the fiducial continuum detection of Elias 27 from the DSHARP data release [77], produced using a self-calibrated image with `CASA`. The image on the right shows the result of applying `RESOLVE` to the Elias 27 continuum ALMA data, displaying the mean sky map. In this case, no self-calibration was applied. The image measures 1500×1500 pixels on the site, corresponding to a field-of-view of 3.33×3.33 mas^2 on the sky. Based on the work of [15–17], we anticipate further improvements in detecting source details when self-calibration is incorporated into the `RESOLVE` framework for ALMA data. The novel procedure will allow the introduction of gain estimations and their uncertainties. In Fig. 16 we show on the top the channel maps detection of $^{12}CO\ J = 2-1$ as provided by [7]. The lower images display two channels at velocities of 1.35 km/s and 2.05 km/s, reconstructed using `RESOLVE`. Each image comprises 2000×2000 pixels, representing an area of 10×10 mas^2 on the sky. The detection of the diffuse emission occurs separating it from the continuum emission. The diffuse emission shapes can be well recovered by imaging each channel separately. Although the combined cube images are not yet as good as the clean reconstruction the preliminary results look promising; there is hope that suitable tuning of hyperparameters allow the reconstruction of good cube images [13]. We expect better images with improvements in code efficiency. In this regard, a combined application with `RESOLVE` and `DeepFocus` seems promising.

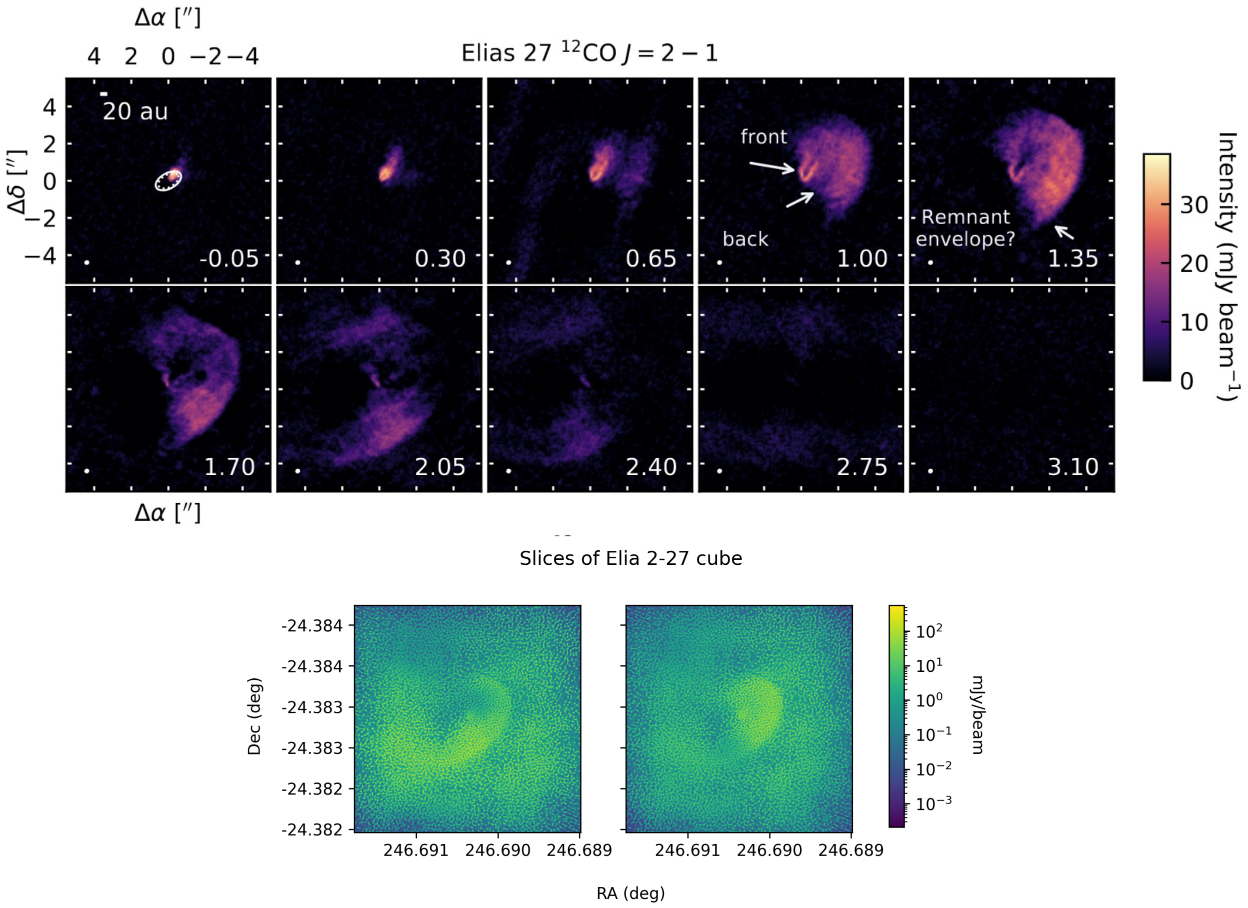

Figure 16: Elias 27 from the DSHARP ALMA project, channel maps of the ^{12}CO $J = 2 - 1$. Top: the fiducial image as given by the DSHARP team [77]. Bottom: `RESOLVE` *mean sky map of the channels close at the LSRK velocity of 2.05 (left) and 1.35 (right) km s^{-1}.This is a proof-of-concept. See text for more details.*

In collaboration with DSHARP team members, we are considering to extract physical information from continuum and cube images with `RESOLVE`. IFT will be employed to develop physics-informed imaging and spectral reconstruction methods, offering comprehensive uncertainty quantification for all scientific results. This approach will provide the significance of any detected continuum emission structures, along with atomic and molecular lines from their disks, outflows, and planets, ensuring a robust scientific interpretation.

4.3.2 Computational aspects

Following the work in [76], Section 5.3, `RESOLVE` was initially found to be 60 times slower than multi-scale `CLEAN`. However, subsequent advancements, as discussed in [13], have significantly improved its performance, achieving speeds up to $140\times$ faster compared to its initial applications to DSHARP data (benchmarking was performed on an older version of `RESOLVE`, as reported in [76]).

For DSHARP imaging, reconstructions using `RESOLVE` required approximately 3–4 hours for 40 iterations of a Maximum-A-Posteriori (MAP) run. MAP provides a faster approximation of the full posterior, assuming the posterior follows Laplace's approximation. In contrast, a MGVI run for the same dataset typically takes around 3 days for 40 iterations. These timings assume the use of 8–16 cores on a cluster to reconstruct a 2000×2000 pixel image across 16–60 frequency channels.

The earlier DSHARP imaging tests using `RESOLVE` were completed in a few hours, though convergence was not accounted for in those runs. We have not directly compared `RESOLVE`'s computational time with `tCLEAN`. A proper comparison would require reprocessing the data using the latest `CASA` version. Such a comparison remains an important goal, especially as `tCLEAN` evolves to include options for Multi-Frequency Synthesis (MFS) and Multi-Term MFS (MTMFS) deconvolution, along with GPU-enabled gridding options, hopefully, for ALMA data.

4.4 Application to RX_J1347.5-1145, Detecting the Sunyaev-Zel'dovich Effect

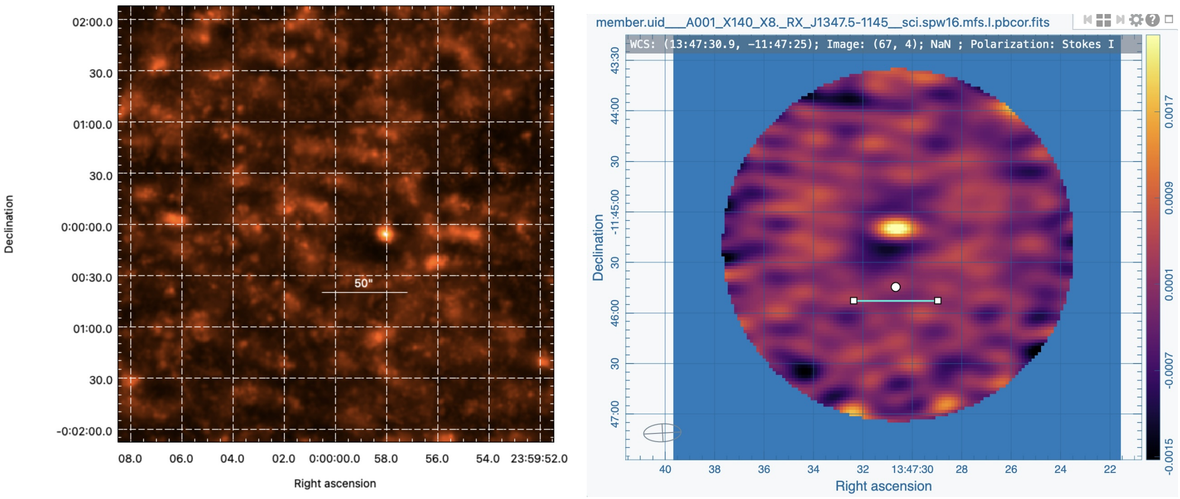

Figure 17: Proof of concept of `RESOLVE` *capabilities to detect the SZ effect in ALMA data. The archived ALMA source* $RX\_J1347.5-1145$*, from project* $2013.1.00246.S$*, scheduling block* $RX\_J1347\_a\_03\_7M$*, mosaic, has been accounted. Left:* `RESOLVE` *mean sky map. Right: The corresponding archived ALMA product, primary beam corrected. A bar of* 50 *arcsec length has been added to both images to compare the detected depths. See text for more details.*

The archived ALMA target $RX\_J1347.5-1145$, from project code $2013.1.00246.S$, was utilized to evaluate the performance of `RESOLVE` in detecting "decrement" like features, such as those caused by the Sunyaev-Zel'dovich (SZ) effect. The SZ effect is due to scattering, which preserves photon number, of the Cosmic Microwave Background (CMB) photons to higher energies, leaving a decrement at longer wavelengths and producing an increment at shorter ones. A comprehensive discussion on the astrophysical significance of SZ effect detection can be found at [78, 79].

For this analysis, a single pointing from the mosaic data in the scheduling block (SB) $RX\_J1347\_a\_03\_7M$ was selected. In Fig. 17, the left panel shows the output of the `RESOLVE` algorithm, applied to a single execution block $(uid://A002/X966cea/X22ee)$. The faint distortion of the CMB, observable at (sub)millimeter wavelengths, appears as a dark region in the image. A region of 50 arcseconds is outlined to provide a sense of the scale of the detected area. The panel on the right shows the archived ALMA product obtained on the same science target and same frequency range. This `tCLEAN`-*ed* image is applied on the whole mosaic, accounting for all QA0 PASS execution blocks, and primary beam corrected. Also in this case a bar of 50 arcsec length has been added for comparison.

The `RESOLVE` application has been recorded in GitLab and accessible clicking at this link. Although we provide an initial proof-of-concepts, we aim at extending the applications of `RESOLVE` to the following ALMA archived sample:

- SPT-CLJ2031-4037: project code $2019.2.00067.S$, MOUS uid: $uid : //A001/X14c3/X326$, SB name: $SPT-CLJ2\_a\_03\_7M$. The aggregate continuum shows a nice SZ detection. The spectral window (SPW) 16 is the one with the most prominent continuum detection of the SZ effect.
- SPT-CLJ2106-5844: project code $2016.1.01175.S$. This is composed by 2 SBs with less prominent decrement features.
- SPT-CLJ2106-5844: project code $2017.1.01649.S$, composed by 2 SBs taken with the 12-m array with less prominent decrement features.
- SPT-CLJ2106-5844: project code $2021.1.01212.S$ taken with the 12-m array. The single SB shows some weak line emission detection in SPW 23 (expected CO line), the continuum may detect the SZ effect although this is not pronounced. The mfs SPW with more pronounced decrement is the 21.
- SPT-CLJ2106-5844: project code $2022.1.00333.S$, MOUS uid: $uid : //A001/X2d1f/X80c$, SB $SPT - CL\_J\_a\_03\_TM1$. The CO line is nicely detected in cube SPW 17, and the aggregate continuum is detected in emission only.

This intriguing sample still presents several unresolved questions. For example, by combining all the Band 3 data from project $2017.1.01649.S$, the study by [80] achieved a detection significance greater than 20 sigma using a Bayesian forward modeling approach for the cluster SPT-CLJ2106-5844. A comprehensive analysis of all datasets for this science target, focusing solely on continuum detection, is expected to yield a highly detailed SZ map (private communication with TM).

Furthermore, the science target SPT-CLJ2031-4037, observed in mosaic mode, can be fully analyzed with the current capabilities of `RESOLVE`. By combining all execution blocks, the resulting image is anticipated to be equally impressive. The continuum-cleaned map should reveal a shock feature that aligns well with the Chandra X-ray observations, lending credibility to the authenticity of this structure.

4.5 Mosaicking and Group OUS Imaging

By adopting a Bayesian approach to data integration, the `RESOLVE` algorithm can be enhanced to effectively combine interferometric and single-dish data. This is achieved by treating each set of data as distinct likelihood functions. Consequently, the combined likelihood is formulated as:

$$\mathcal{P}(d|\vec{s}) = \prod_{pq} \mathcal{P}^{p}_{\mathrm{SD}}(d|\vec{s}) \cdot \mathcal{P}^{q}_{\mathrm{IF}}(d|\vec{s}), \tag{1}$$

where p and q represent the single-dish (SD) and interferometer (IF) data, respectively. This representation includes the condition that when imaging celestial objects larger than the primary beam size of a radio telescope, it is essential for both single-dish and interferometer instruments to map the science target in order to capture the object's full extent. The observations, i.e. the associated data d, jointly constrain a common sky model $\vec{s}$. This model is reconstructed by incorporating a prior PDF and applying advanced optimization techniques (refer to [71, 72]).

To demonstrate the efficacy of this approach on mosaicking, we present reconstructions of synthetic data and compare them to those obtained using the standard `tCLEAN` method (see [12]). The synthetic data, derived from an image of M51 (depicted in the upper left panel of Fig. 18), was generated using the `simalma` task in `CASA`. The results indicate that the `RESOLVE`-based algorithm achieves satisfactory reconstructions, with a comparative analysis of residuals revealing that `RESOLVE` generally provides higher fidelity in capturing fine details of the input data. This is evident from the reduced spread of residual pixels (see Fig. 18). These findings suggest that `RESOLVE` can deliver accurate reconstructions, provided that uncertainty quantification from the calibration is understood and properly propagated to imaging processes.

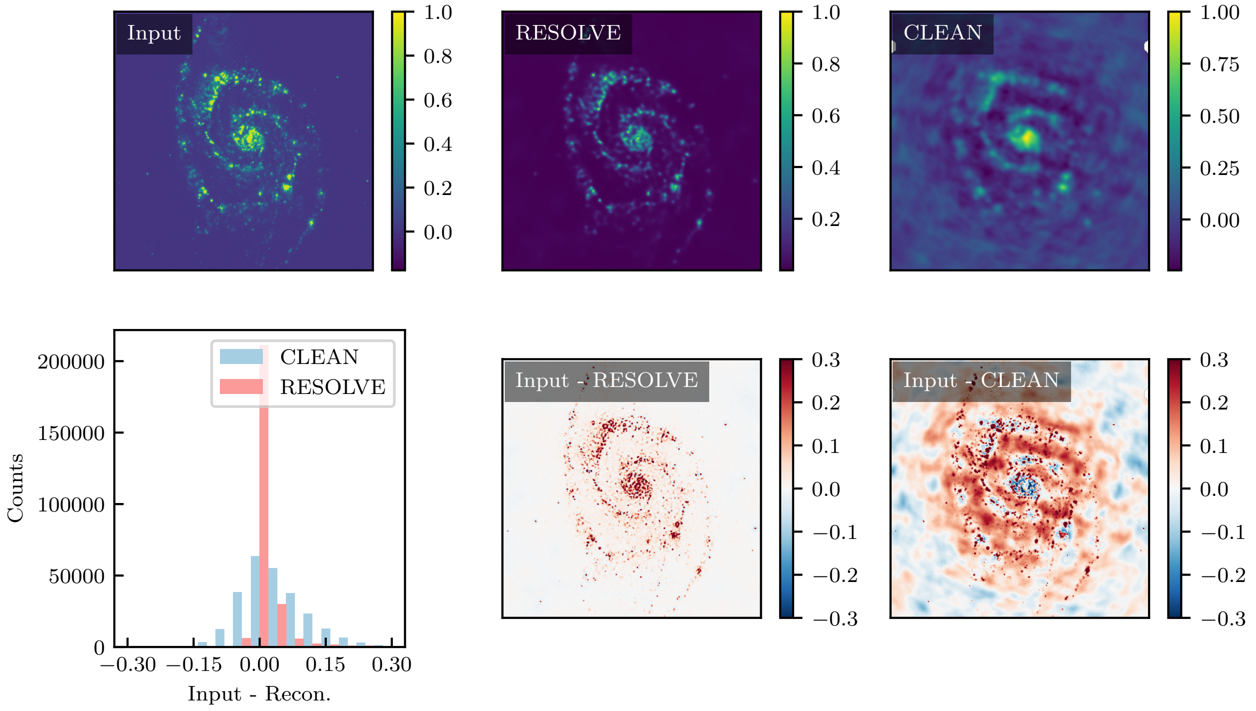

Figure 18: Comparative analysis of `RESOLVE` and `CASA tCLEAN` on synthetic multi-pointing ALMA data. The top row displays, from left to right, the input image of M51 used for generating the synthetic ALMA data, and the reconstructions of M51 obtained using `RESOLVE` and `CASA`, respectively. The bottom row presents the residuals relative to the input image. The left panel shows binned image pixel residuals comparing `RESOLVE` (red) and `CASA` reconstructions to the input, while the middle and right bottom panels display the image residuals on the image plane.

Our next step is to apply `RESOLVE` to perform the GOUS imaging on ALCHEMI data. Comparisons with other techniques, as Feathering [81] and the joint convolution algorithm [82], shall be performed.

4.6 Ongoing developments and Outlook

The technique `RESOLVE`, applied to ALMA data, provides for image reconstruction from sparse samples and proper uncertainty quantification. `RESOLVE` is based on IFT, which extends Bayesian probability theory to physical fields. Bayesian probability theory allows one to extend from point estimates to a distribution of solutions which enables the derivation of optimal values and uncertainties for the quantities of interest. Input information to the imaging technique is the observed and calibrated visibilities. The dirty beam is derived from the uv coordinates and incorporated into the Response operator in the deconvolution process. The input data are processed within a probabilistic approach allowing one to model the celestial signal and noise in the data in one unique algorithm. Products of the technique are reconstructed ALMA deconvolved image and its uncertainty map, the power-spectrum and its uncertainty, estimates of input parameters and their uncertainties. Through simulated and real data, we demonstrate that `RESOLVE` is applicable to ALMA data. Detections of extended emissions and structures in protoplanetary disks are promising. The application to Elias 27 for continuum de-

tection, when compared with self-calibrated `tCLEAN`' solution, `RESOLVE` is capable to improve the detection of the science target and the image of the diffuse emission.
The `RESOLVE` algorithm has been applied to detect the Sunyaev-Zel'dovich (SZ) effect (decrement detection) and to analyse mosaics of images.
Substantial efforts have been made by the IFT team to make the algorithm computationally efficient. Variational inference has been introduced with the works of [71] and [72] making feasible to have estimates on high-dimensional posterior distribution functions. JAX binding [73] is making possible to use CPUs and GPUs. A new version of `RESOLVE` [13] is now at least 140 times faster than the previous version when executed on a GPU. A systematic comparison between the runtime of `RESOLVE` and `tCLEAN` has not been performed. However in some applications the runtime of `RESOLVE` and `fast-resolve` was measured. Based on results VLA data Cygnus A reconstructions of `RESOLVE` is approximately 60 times slower than `wsCLEAN` and `fast-resolve` comparably fast when executed on GPU. Nevertheless on other datasets these timings might be different, and also the runtime of `wsCLEAN` and `CASA tCLEAN` might be different. The exploration of cube imaging required intense work, but a proof-of-concept is given. IFT is allowing a joint continuum and line detection. However, the models have to be refined in order to provide cleaner images. At the time of the writing, `RESOLVE` is tested on our ARC cluster at ESO. We will provide first assessments on the execution speed and define its computational cost.
Furthermore, we want to exploit the unique capability of the Bayesian analysis to (1) combine the information of different instrument types, as ALMA and JWST; (2) combine archival data as performed by [35]; (3) create dust maps of deep observations as performed by [83].
Efforts to extend the capabilities of the `RESOLVE` algorithm must also address user concerns regarding the interpretation of super-resolution images. The weighting scheme utilized by `RESOLVE` is adaptive and based on baseline noise, operating as a natural weighting scheme. While the output resolution is a free parameter determined by the model, it is constrained by the chosen pixel number and field of view, which together define the pixel size (or pixel resolution). Importantly, the maximum achievable resolution is set by the largest baseline in the data. Therefore, selecting an appropriate input resolution consistent with the data's intrinsic limitations is essential. Comparative analyses of `RESOLVE` with `tCLEAN` under specific weighting schemes (e.g., as outlined in [76]) are under consideration to provide further insights into performance differences.
To improve usability and instill greater confidence in the outputs, several future developments are proposed. (1) Provide Angular Resolution Metrics: Introducing a dedicated task to compute and output the effective angular resolution of `RESOLVE` images will enable users to directly compare the effective resolution with the native resolution of the input data, as well as any super-resolution enhancements applied. (2) Expand Uncertainty Mapping: The algorithm already provides uncertainty maps. These can be extended to illustrate spatial variations in resolution, offering users additional context on image reliability. These outputs, in FITS format, shall include headers specifying the achieved resolution. (3) Write a User Guide for Super-Resolution: A comprehensive user guide will be developed to educate users on interpreting super-resolution results. This document will cover the algorithm's methodology, limitations, and best practices for evaluating outputs to ensure informed usage.

These advancements will not only expand the algorithm's functionality but also address key challenges in user interpretation, thereby facilitating broader adoption and more effective application of `RESOLVE` in astronomical research.

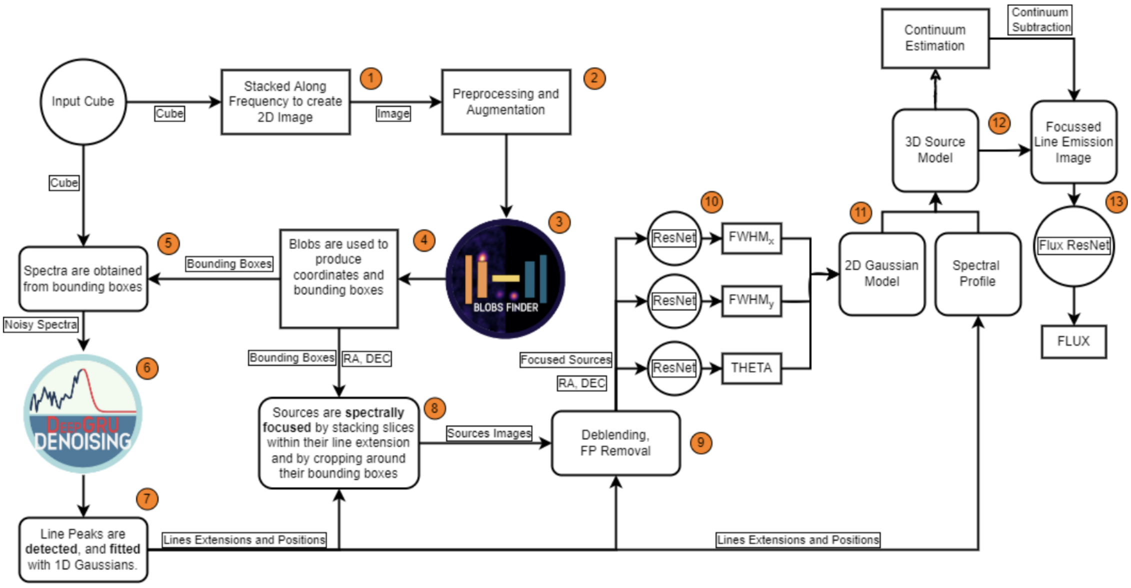

Figure 19: *Deep Learning Pipeline schema. Numbers indicate the logical flow of the data within the pipeline.*

5 DeepFocus

The Deep Learning Pipeline, named `DeepFocus`, is fully described in the refereed article accepted for publication on MNRAS: click on this link to be redirected to the article [2].
`DeepFocus` is designed as a modular system allowing to expand the pipeline according to the problem at hand. In operations we always need to cope with new modes or instrumental improvements (e.g. electronic upgrades, as [3]). New units can be connected to the pipeline as well as replacements with others during software refinements. The schema of the composition of the Pipeline can be seen in Fig. 19. Currently `DeepFocus` is composed by six Deep Learning models: a Convolutional Autoencoder for source detection within the spatial domain of the integrated data cubes (Blobs Finder, stage 3), a Recurrent Neural Network for denoising and peak detection within the frequency domain (DeepGRU, stage 6), and four Residual Neural Networks for source characterization (ResNets, stage 10).
The algorithm solves for $I^D = I_{db} * I + n$, where I^D, I_{db}, I and n are the ALMA dirty cube, the dirty beam, the target sky image and any additional noise in the image, respectively. The information contained in each channel in the ALMA cube is accounted in the whole frequency domain, correlating information between pixels along the frequency axis of the cube. It has been demonstrated that the combination of spatial and frequency information has the capability to improve completeness while decreasing spurious signal detection. Comparison with `tCLEAN` perfomed analysing a set of 1000 simulated ALMA data cubes indicated that `DeepFocus` improves in speed up procedures by a factor of at least 140.

5.1 Proof of concepts

Supervised ML techniques need training, testing and validation to become that powerful tool to outperform any current algorithm to date in processing time. Most of ALMA archived data are galaxies, with compact and almost point-like shape. The potential capabilities of this algorithm have been tested on ALMA SV interferometric data: BR1202-0725 [84]. Please refer to the article [25] (reachable clicking this link) to get more details on this application (Subsection 2.2 and Figure 4).

This trial shows that both galaxies in the ALMA cube are detected with `DeepFocus`, while the standard approach in `CASA` missed the serendipitous quasar (southernmost object). In [84], self-calibration was performed to detect this serendipitous source. The source fluxes derived by `DeepFocus` agree with the ones reported in [84]. The computing time for the image restoration of one ALMA cube for BR1202-0725 with `DeepFocus` occurred in $\sim$ 35 μs.

5.2 Simulations

Simulated ALMA cubes are used for training/testing/validation as well as reliability and quality assessment of the developed Deep Learning pipeline. Section 3 of the journal article “3D Detection and Characterisation of ALMA Sources through Deep Learning” [2] provides a detailed description of the developed artificial data sets. The `CASA` simulator capabilities [20], [85] were extended for the creation of ALMA cubes. 2D Gaussian Components in the spatial plane with 1D Gaussian component (emission lines) in frequency space are used to create the artificial emission lines in addition to the continuum signal. ALMA Cycle 9 C-3 configuration with 43 antennas was chosen within the `simalma` task in `CASA`. ALMA interferometric MSs are created employing the `CASA`’ `simobserve` task. Dirty cubes are produced employing `tCLEAN`. Corrupted by white noise, the dirty cubes’ RMS is adjusted to a wished SNR.

A set of 5000 ALMA cubes (360x360 pixels, 128 channels and total bandwidth of 1.28 GHz) are created with a source at the center and other randomly distributed in the image with random extension. The brightest source is located at the center and characterized by a SNR $>$ 10. The minimum and maximum flux densities generated are respectively 0.97 and 407.4 mJy/beam. Uniformity on the distribution of simulated source parameters is achieved (Fig. 3 in [2]). See upper and middle rows of each set in Figs. 26 and 27 for some examples of produced ALMA dirty cubes and model images.

In these simulations, only limited test cases were given as training set, mainly consisting of continuum and single to multiple Gaussian shape lines. Widening the training set to complex emissions as well as empty sky images was beyond the scope of this study.

5.3 Novelty of the method

The pipeline architecture can be roughly divided into three phases, based on the assumptions made on the data. First, sources are assumed to be present within the image and the algorithm is trained to detect those sources. Second, sources to be deemed true must show emission lines in the frequency domain. The algorithm searches for those emission lines by removing noise in order to boost SNR and recognise spectral peaks. If spectral peaks are found, the algorithm preserves the sources for characterization, otherwise those initial identifications are discarded as false detection. The latter operation is performed through SNR and geometrical criteria involving the reference integrated dirty image and all images produced by integration along the detected emission ranges. Third, sources passing the selection criteria are fed to an array of ResNets to regress the morphological source parameters.

As future outlook of the development of `DeepFocus` and its applications, we plan to implement a cross-check using training exclusively on empty images containing only noise. This approach aims to assess the reliability of the algorithm by testing its performance in detecting spurious signals in purely noisy fields.

Currently, these tests on empty images have not yet been conducted, as the training phase for ML algorithms on such data was not considered. We want to incorporate these tests as future development, enabling a more comprehensive validation of the algorithm’s detection capabilities. This step will provide further confidence in the robustness of `DeepFocus`’s results, particularly in scenarios with low signal-to-noise ratios.

5.3.1 Blobs Finder

Blobs Finder solves for the deconvolution problem in the image domain, i.e. to recover the normalized denoised integrated sky images from the integrated dirty images (stage 3 in Fig. 19). The output probability maps are censored and all pixels with probability higher than a given threshold are connected into potential sources through a friend of friend algorithm. Bounding boxes are extracted around the islands of connected pixels to define source spatial boundaries. Bounding boxes around source candidates are used to extract dirty spectra from the input dirty cube by summing, for each frequency slice (channel), all pixels within the bounding boxes. Figures 20, 21, and 22 show, respectively, an example of an input integrated dirty cube containing 6 simulated sources (green boxes) with two spatially blended sources, the target sky model image (green and red highlight the target and the predicted bounding boxes), and the 2D prediction map (red bounding boxes).

5.3.2 DeepGRU

The obtained spectra are standardized and fed to DeepGRU (step 6 in Fig. 19) which is tasked to solve a 1D denoising problem and outputs 1D probabilistic maps of source emission lines which are then analyzed in search for peaks. Each peak is fitted with a 1D Gaussian function. Position z and extension $\Delta_z = 2 * FWHM_z$ (where $FWHM_z$ is the FWHM of the Gaussian peak) are recorded. Fig. 23 shows the dirty spectrum extracted from the two blended sources shown in Fig. 22, and the DeepGRU's predicted emission probability map. Blue and red vertical bars limit the true and predicted, respectively, emission ranges of the two sources within the spectrum. In order to detect possible false positives produced by Blobs Finder, all potential candidates showing no meaningful peak in their spectra are removed. If more than one peak is found alongside the spectrum, three possibilities may arise: detected peaks may indicate distinct celestial sources which are spatially blended, detected peaks may belong to the same source, or one or more peaks are false detection(s). None of these possibilities can be excluded a priori. On each peak we perform *spectral focusing*, i.e. we crop a 64×64 pixel image around the source center in the spatial plane and integrate within the peak extension in frequency. In order to estimate the SNR of a source, two SNR measurements are accounted:

Global SNR:

$$SNR \overset{\text{def}}{=} \frac{median(x_s(r))}{var(x_n(R-r))} \tag{2}$$

where $x_s(r)$ are the pixel values of the source within a radius r inscribing the bounding box, and $x_n(R-r)$ are the pixel values within an annulus of internal radius r and external radius R which has the same area of the inscribed circumference;

Pixel SNR:

$$snr \overset{\text{def}}{=} \frac{x_i}{var(X)} \tag{3}$$

where x_i is the value of the given pixel, and $var(X)$ is the variance computed on the full image.

The two SNR ratio measurements are used to distinguish falsely detected from true sources and to deblend overlapping sources within a blob. Fig. 24 summarises the false positive detection pipeline. In case of not blended detections, the process works as follow:

- if $SNR \geq$ 6 eq. (2) (empirical bright source SNR threshold) in the integrated dirty cube and the source is not flagged for deblending, the detected source is focused;
- if $SNR <$ 6 eq. (2) in the integrated dirty cube, the detected source goes through a check for focus:
  - if SNR eq. (2) increases, the source is kept;
  - otherwise the source is discarded as false positive.

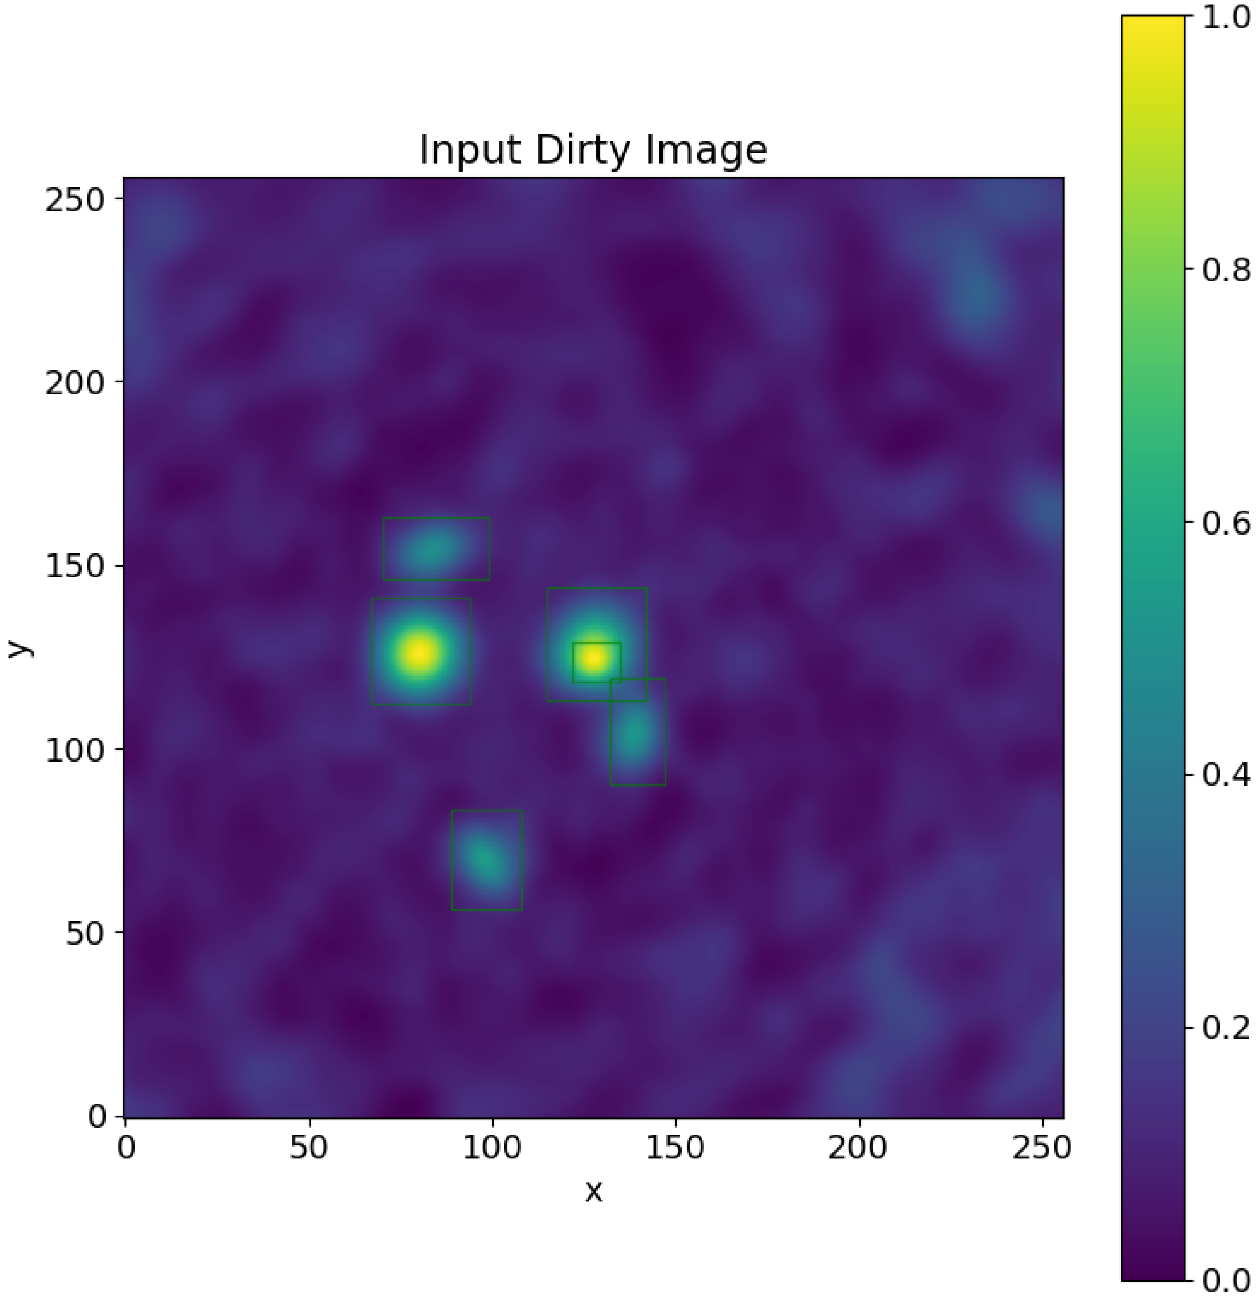

Figure 20: *An example of Blobs Finder's input 2D integrated dirty cube produced by integrating an input dirty cube over the entire frequency range. Superimposed in green, are the target bounding boxes outlining the emissions of the* 6 *sources present in the cube. The image contains an example of two spatially blended sources located around the centre of the image, one is a bright point-like source, the other a fainter and diffuse source laying behind.*

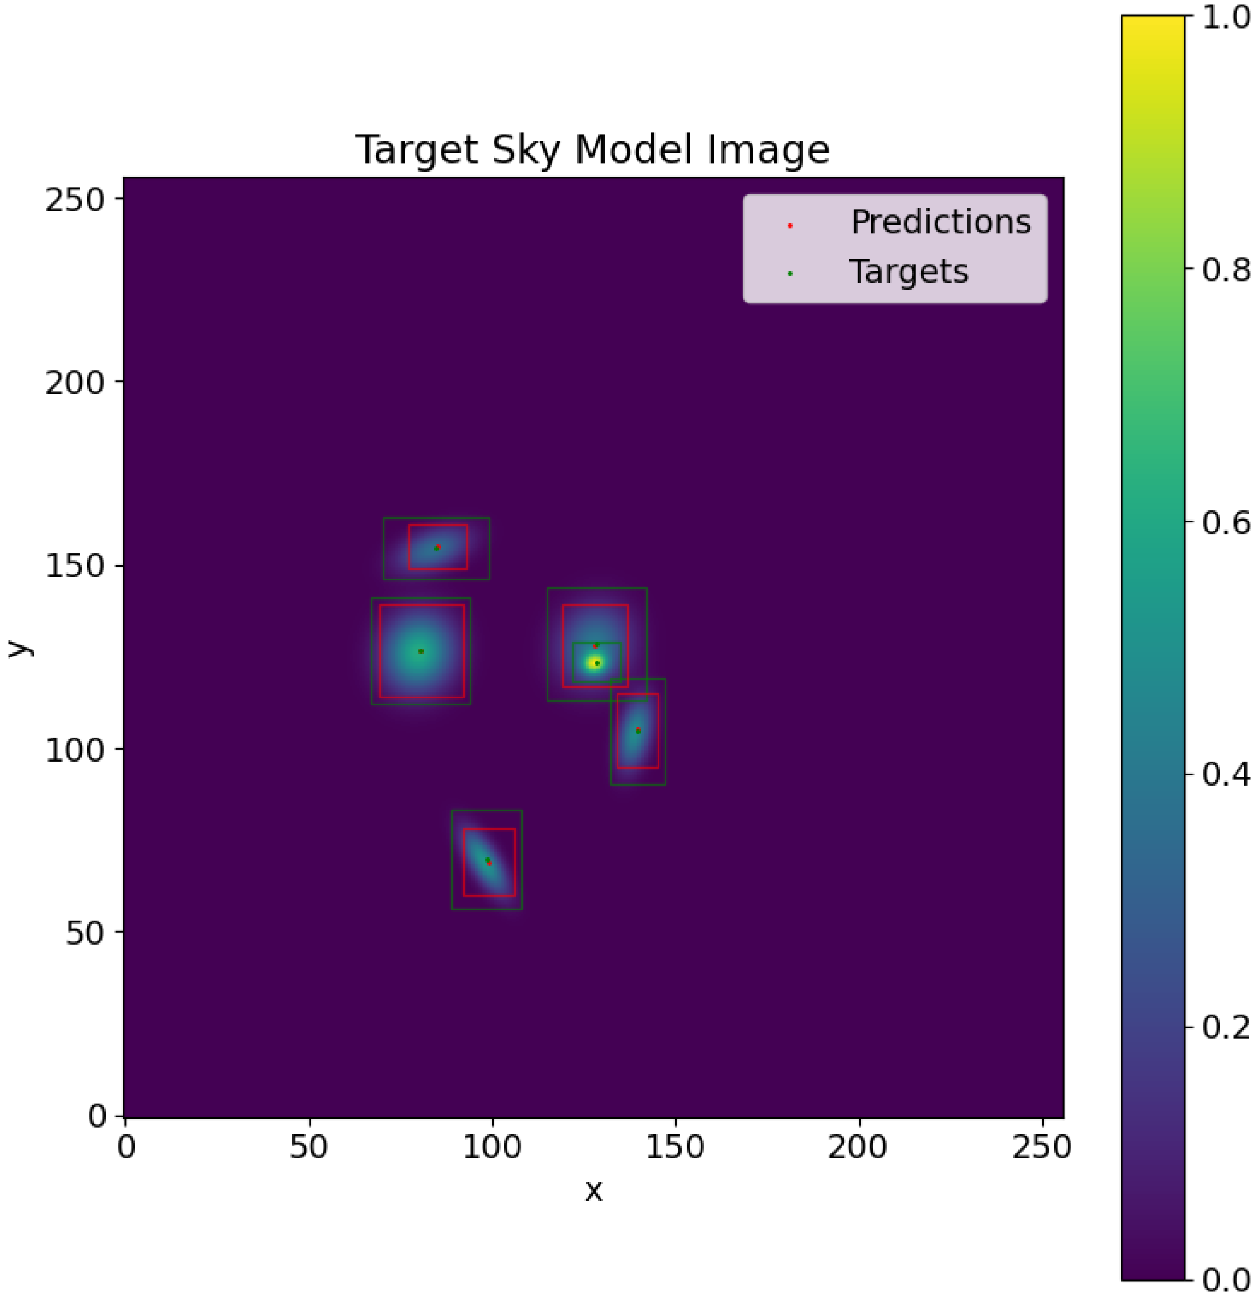

Figure 21: An example of Blobs Finder's target 2D Sky Model image with the target bounding boxes highlighted in green and the predicted bounding boxes extracted through the tresholding operation on Blobs Finder's probabilistic output, highlighted in red. Predicted and true bounding box centers are also plotted as, respectively, red and green dots.

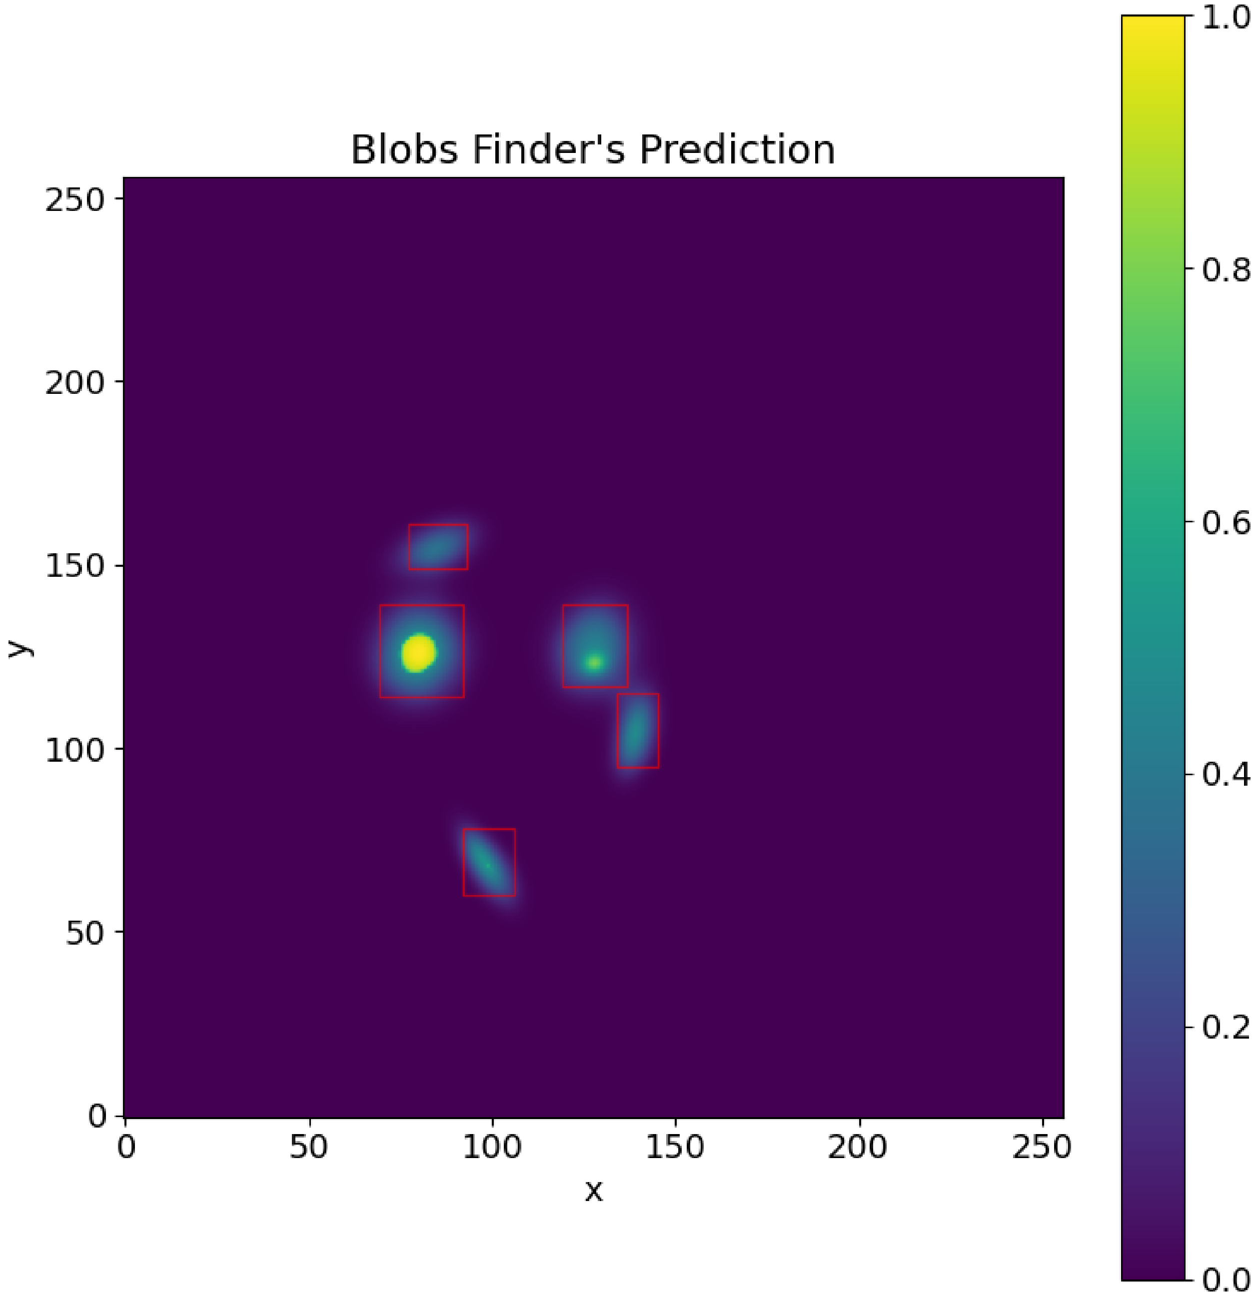

Figure 22: An example of Blobs Finder's output 2D probabilistic source detection map with the predicted bounding boxes extracted through thresholding, highlighted in red.

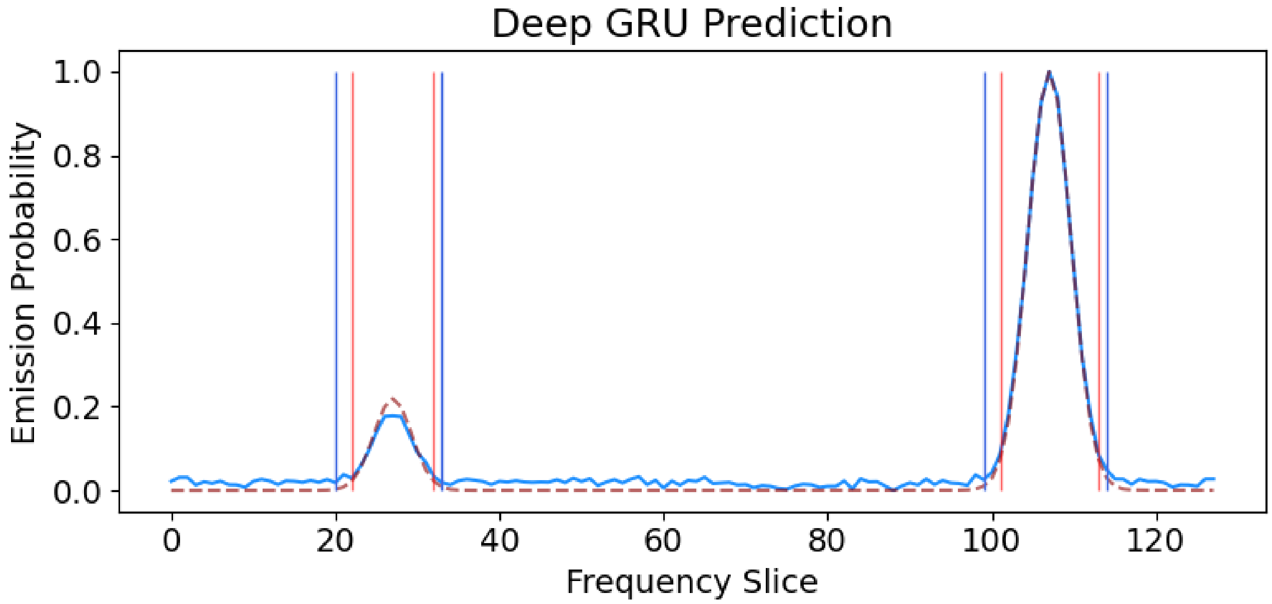

Figure 23: *In blue the dirty spectrum extracted from the central source bounding box predicted by Blobs Finder (Fig. 21), in dotted-red the DeepGRU's prediction. Vertical blue bars delimit the true emission ranges, while red bars the predicted emission ranges. A secondary fainter source emission peak is detected by DeepGRU and thus the source is flagged for deblending.*

This is the condition marked with 1 in Fig. 24. The latter case, in fact, could only happen if the source is integrated outside its true emission peak, for example over a noise spike. If more than one peak is found in the potential source spectrum, then there is a chance that multiple blended sources make the blob (detected by Blobs Finder). In case of blended sources, the procedure for the identification of true sources proceed as follows:

1. focusing on the highest peak (primary peak) by integrating within its extension allows for the SNR eq. (2) calculation. The same logic described above is followed.
2. the snr measurement eq. (3) is used to identify the pixel with strongest intensity in the image $p(x, y)$. This reference pixel is used in the next phase of the deblending process.
3. secondary peaks not overlapping in frequency with the primary peak are analysed:
    - the snr measurements eq. (3) find the reference pixel in the image $s(x, y)$.
    - friend of friends algorithm is used to link pixels around the new reference pixel in $s(x, y)$ until a saturation level is reached when calculating the SNR eq. (2) iteratively.
    - a bounding box is created to encompass all the selected pixels, and a $[64, 64]$ pixel image is cropped around the bounding box.
4. secondary peaks overlapping in frequency with the primary peak are inspected:
    - if the primary and secondary peaks coincide spatially ($p(x, y) = s(x, y)$), then the secondary peak is discarded as a false detection (condition marked as 2 in Fig. 24).

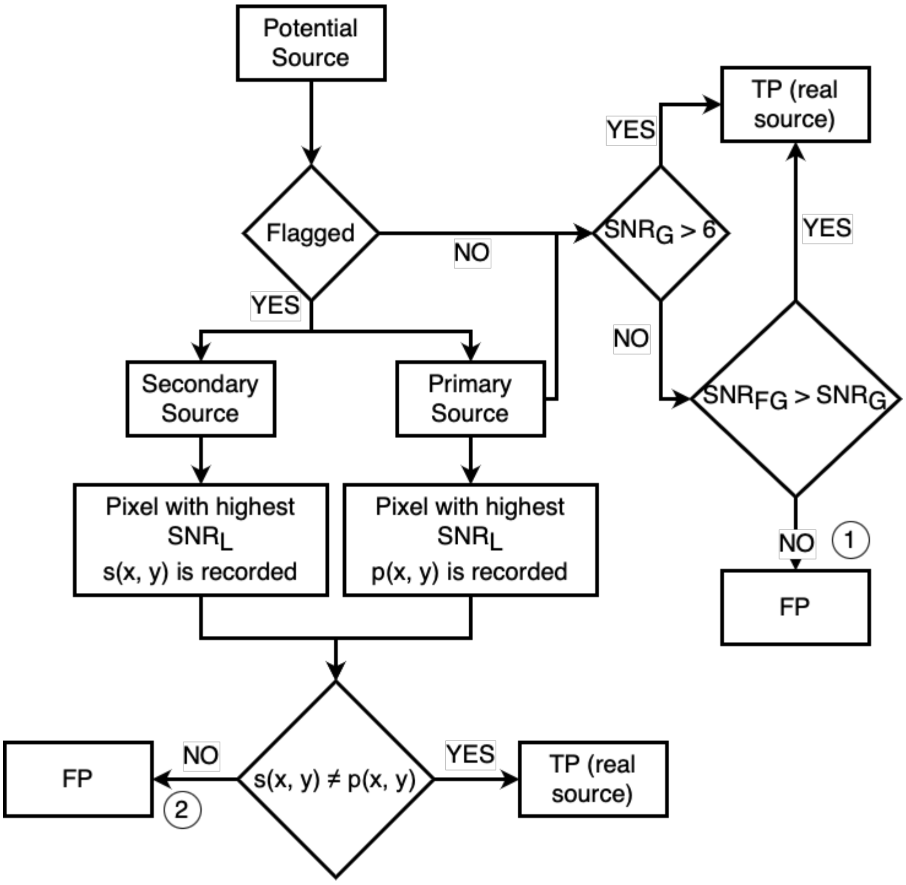

Figure 24: Schema of the False Positives detection and source deblending pipeline which constitutes step 9 *in Fig. 19. Numbers* 1 *and* 2 *show two possible conditions: a potential source being defined as false positive and discarded from further analysis. The subscript* FG *(focused global) indicates that the Global SNR is measured on the focused source, while* L *implies a (local) Pixel SNR measurement. Flagged expresses that multiple peaks are detected within the potential source's spectrum and thus the source is flagged for deblending.*

  - e.g. DeepGRU may predict a single peak as two separate peaks or Blobs Finder predicts a single true source as two very close blobs.
- if $p(x, y) = s(x, y)$ but SNR eq. (2) increases : the secondary peak is deemed as part of the primary source and the source emission range is extended accordingly.
  - it may happen if DeepGRU overpredicts the true emission range.

Finally all spectrally focused sources with SNR lower than 1 eq. (2) are flagged and removed.

Fig. 25 gives an example of Spectral Focusing applied to the potential sources detected by Blobs Finder and DeepGRU in the test cube already displayed in Fig. 20. By focusing on the two peaks detected by the DeepGRU Fig. 23, the two blended sources produce two different images (Focused Source 0 and 1) which can be analysed independently.
An additional advantage of focusing is the improved dynamic range of the detected signal with respect to the registered one in the reference dirty integrated image.

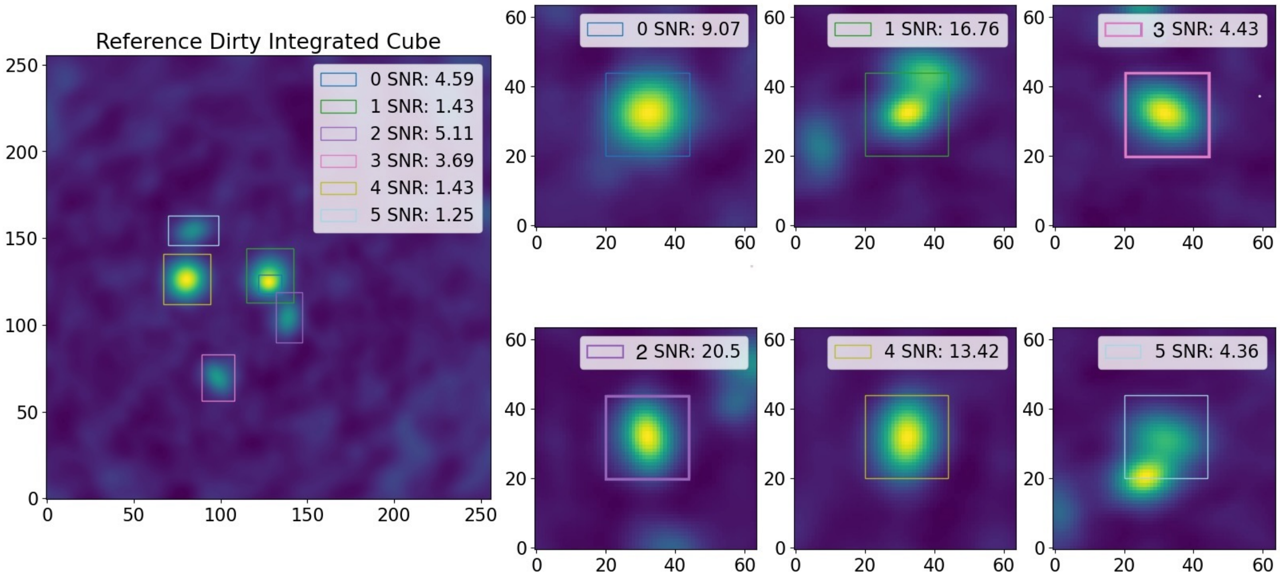

Figure 25: *An example of source spectral focusing of sources within a test set image. On the Left, as reference, we plot the dirty integrated cube with the predicted 2D bounding boxes obtained by Blobs Finder highlighted in different colours. The legend matches the source number to the bounding box colour in the image and the measured Global SNR eq. (2). On the right, there are the 6 Spectrally Focused images obtained by integrating over the predicted line extensions found by DeepGRU and cropping a $[64, 64]$ pixel image around Blobs Finder's predicted bounding boxes centres. In each focused image it is also showcased the measured Global SNR. A substantial increase in SNR occurs when sources are focused around their actual emission ranges.*

Generalization to more complex spectral profiles: DeepGRU makes the assumption that emission lines are mostly consisting of a single Gaussian. This simple approximation is valid for a large fraction of ALMA targets, but there are other sources with complex velocity structures such as lopsided gas distributions in galaxies, combinations of Giant Molecular Clouds, etc. In order to generalize the treatment to cover more complex spectral profiles, improved simulations are needed to train the model. DL models often struggle to generalize beyond the scope of their training datasets. To achieve robust detection of more complex emissions, our primary goal is to enhance the quality of the simulations. This, however, demands significant effort and a substantial investment of human power. The improvement of the simulation has been planned as follow: 1. Employment of multiple ALMA array configurations and observational parameters to address QSO simulations and, e.g. through the collaboration with the SKA Source Finding Focus group, complex morphologies of elliptical and spiral galaxies are generated through physics based modelling. Other objects with complex structures (for instance gas distributions, molecular clouds) and multiple spectral lines are accounted in a second step in the simulations; 2. Recover a targeted and comprehensive selection of simulated ALMA observations of interesting objects. Real data and simulated data will be labelled as True and False and a Generative Adversarial Network will be tasked to modify the simulated images until they cannot be discerned anymore from the real data. This will allow us to generate truly realistic moc data on which to re-train our pipeline. See Section 6 for more details.

5.3.3 ResNets

Prediction of morphological parameters of the detected sources (FWHM, coordinates x and y, projection angle pa) is performed by ResNets (stage 10 in Fig. 19). Celestial coordinates are computed as photometric baricenters (pixel-weighted centers) of the Blobs Finder predicted bounding box. Source fitting in spatial and frequency domains are combined to create a 3D Gaussian profile.

A 3D segmentation map is created. The segmentation map is dilated by a factor of 1.5 to account for the convolution process spreading the continuum and the line emission signals in the image. This is performed to make sure that all the source signal is contained within the 3D segmentation map. A dilated 3D segmentation mask is used to create the model-masked cube by multiplying it with the dirty cube. The inverse mask is instead used to capture the continuum cube. The continuum image is created by averaging the continuum cube in frequency. The line emission cube is created through the following formula:

$$L_z[x,y] = M_z[x,y] - f(z) * C[x,y] \quad \text{with } z \in \Delta_z \tag{4}$$

where $L_z[x,y]$ is the 2D line emission image at slice z, $M_z[x,y]$ is the model masked 2D image at slice z, $C[x,y]$ is the continuum image and $f(z)$ is the 1D continuum model. The line emission cube is integrated along the frequency to create the line emission image which is fed to a specialized ResNet predicting the source flux density in mJy/beam.

In summary, ResNets allow us to estimate morphological parameters of the detected sources, measure their continuum and the line emission in addition to create the reconstructed cubes. Please note that in case of overlapping/blended sources, the fitting of the morphological parameters is executed simultaneously. Each detected source (primary or discovered seconday source) is fed to the ResNet to regress the source morphological parameters. The characterization is thus performed simultaneously and no information about the two sources, or any previous source seen in inference, is employed to predict the source parameters.

5.4 Train, Test, Validation

The 5000 simulated ALMA Dirty cubes (described in Subsec. 5.2) are grouped in sets and a split ratio of 60-20-20 to train, test and validate is employed.

- ○ The training set is used to train the DL models within the pipeline.

- The test set is used to measure the pipeline performances in detecting sources and in regressing their parameters.
- The validation set is used to measure the training progress and assess generalization capabilities.

5.4.1 Training and Validation

The training allows the model to initially learn a median representation of the data (which should capture information about the dirty beam and noise patterns), and then refine its understanding to account for nuances in the data, such as source positions and morphological properties (e.g., the shapes and sizes of galaxies). Details on the training strategies can be found in Subsec. 4.5 of [2]. For more information, refer to this link.

At first, training occurs in parallel on the whole pipeline starting from pairs of dirty input images (dirty cubes integrated along the frequency) and target sky model images (target sky model cubes integrated along the frequency) with the Blobs Finder model. DeepGRU is trained on pairs of dirty spectra (extracted from the dirty cubes) and clean spectra (extracted from the sky model cubes). The DeepGRU predictions are used in combination with Blobs Finder's predictions to extract the spectrally focused galaxy images. Targets for training are the simulated source parameters. The three ResNets for morphological parameters estimation are trained simultaneously. In the first training iterations, care is taken in DL model to prevent overfitting. Successively, each model is trained independently on the un-augmented training set predictions of the previous model. In this way, each model corrects for biases introduced by the previous one.
Validation is used to fine tune the training stage. Sec. 4 of [2] provides the description of the validation loss implemented in the algorithm.

5.4.2 Testing: Accuracy evaluation of detected sources after training

The test set is used to evaluate whether the algorithm can generalize well to an unseen dataset (i.e., the 1000 ALMA simulated cubes). The detection capabilities of Blobs Finder and DeepGRU are evaluated by quantifying the overlap between the Ground Truth and the predicted region, using the Intersection over Union (IoU) metric. For Blobs Finder, the 2D IoU between the true 2D bounding box and the predicted one is measured. For DeepGRU, the 1D IoU is measured between the true emission ranges and the detected ones. At least 60% of the 3D emission range of a source must be captured for it to be classified as a true positive (TP). A threshold of 0.6 IoU is chosen to ensure that 90% of the true emission range is captured within the predicted region, given that the line emission image is created using a dilated segmentation mask.

Blobs Finder succeeds with an 89% efficiency and a 0.1% contamination. Spectra from the detected dirty cubes are extracted and fed to DeepGRU, that provides a 99% efficiency and a 0.02% contamination. Sources are "spectrally focused" within the predicted frequency emission ranges Δ_z, and SNR checks are made. It allows to further investigate false (FP) and true positives. The full logic of the FP removal process is shown in Fig. 24. Blobs Finder's false detection is eliminated by DeepGRU, as described in conditions 1 and 2. Figs. 26 and 27 show some examples of Blobs Finder predictions on the test set. For each block, the upper, middle and bottom rows show the input integrated dirty cubes, the target sky models, and Blobs Finder predictions, respectively.

The clean peaks found by DeepGRU are characterized by the ResNets. Fig. 18 of [2] shows the scatter plots of the true parameters versus the predicted ones and the corresponding residuals histograms. The vast majority of residuals lies within $\pm 1\sigma$, indicating that the process is perfectly under control.

Last, `DeepFocus` is shown to obtain better results than other ML techniques when applied on the same test data set. Please refer to Sec. 5 of [2] for more details.

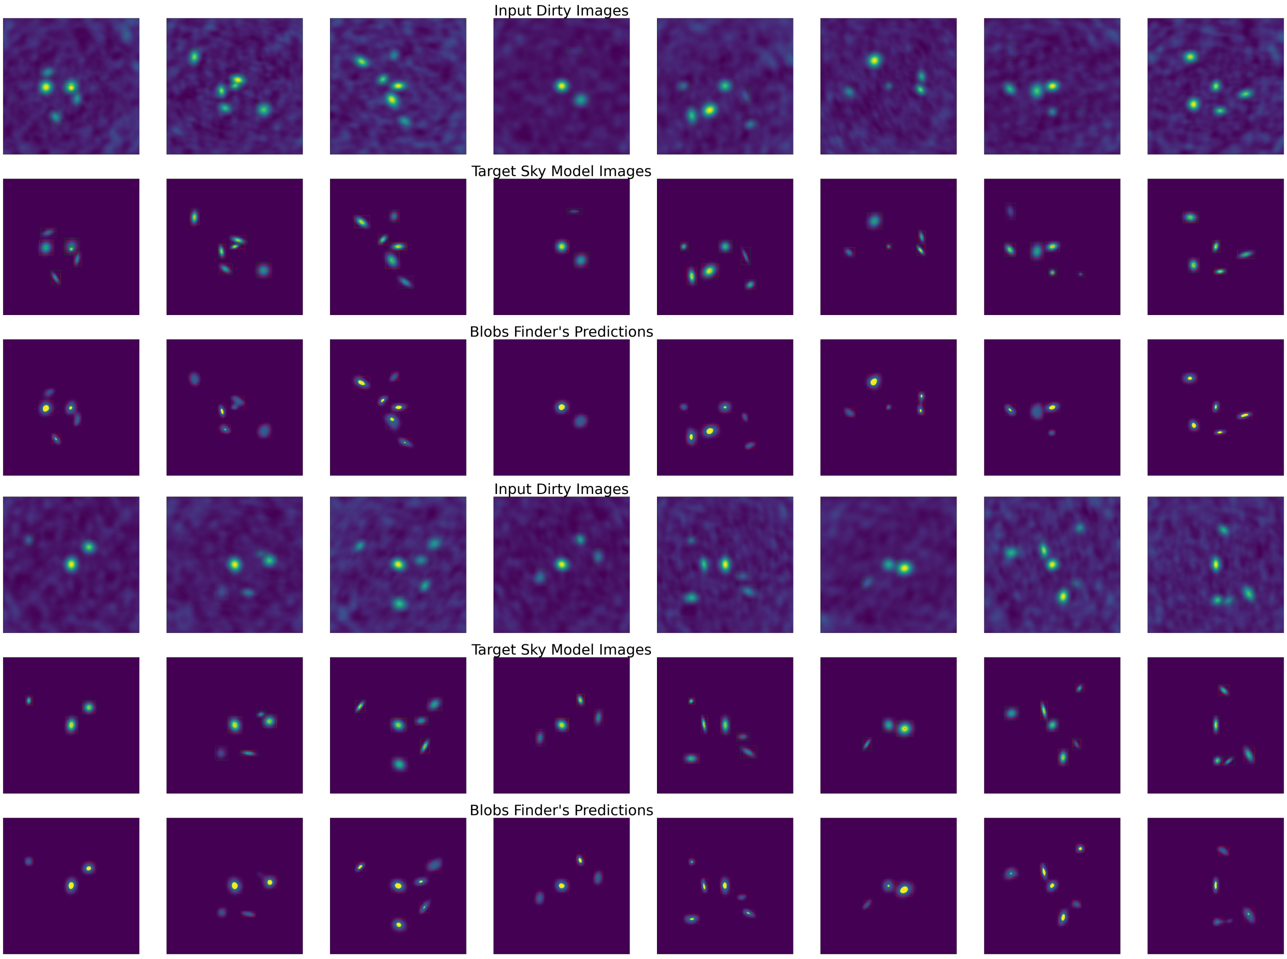

Figure 26: Examples of Blobs Finder predictions on the test set. The first row shows input integrated dirty cubes, the middle row shows the target sky models, and the bottom row shows Blobs Finder predicted 2D Source Probability maps. Green outlines (in the dirty and sky model images) indicate true bounding boxes, while red outlines show predicted bounding boxes extracted by thresholding the probability maps.

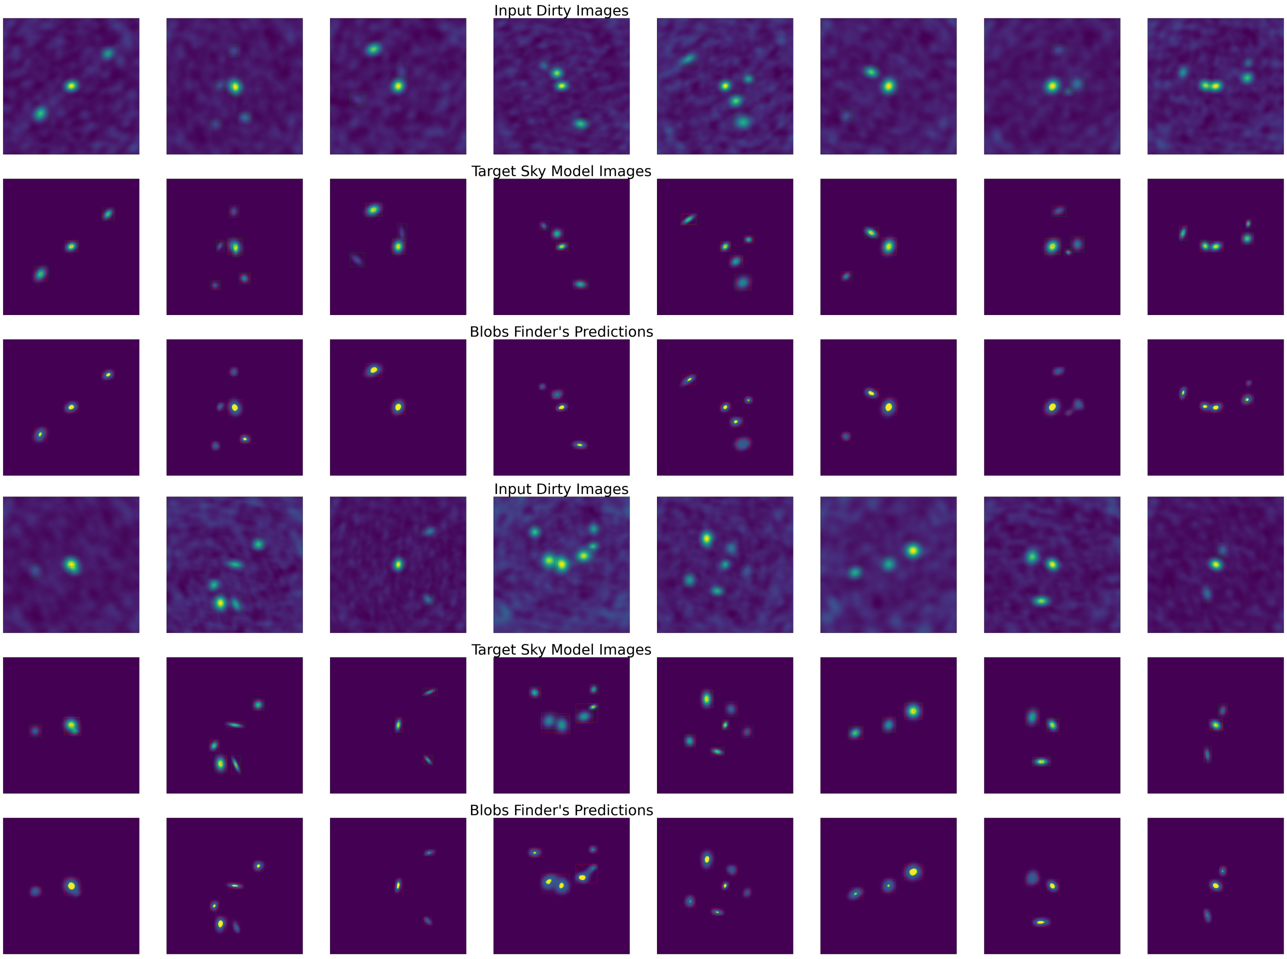

Figure 27: Examples of Blobs Finder predictions on the Test Set. The first row shows input integrated dirty cubes, the middle row the target sky models, and the bottom row, Blobs Finder predicted 2D Source Probability maps. In green are outlined (in the dirty and sky models images) the true bounding boxes, while in red the predicted bounding boxes extracted by thresholding the probability maps.

5.5 Comparison with `tCLEAN` and speed-up estimation

5.5.1 Application of Simulated Data for Comparative Analysis

Preliminary results on comparing the capabilities of `DeepFocus` with respect to `tCLEAN` on a sample of **1000 simulated ALMA cubes** (256x256x128) are provided. In Fig. 28 part of the sample is shown with the first and second columns representing a sample of the integrated ALMA cubes and sky models. The third column shows the solution provided by `DeepFocus`. `tCLEAN` solution run with 200 cleaning iterations is given in the last column. Both algorithms are capable to find all sources in this sample. Residuals are measured to compare the true sky model reconstruction performance of the two algorithms: Blobs Finder's residuals are within $\pm 1\sigma$, while for `tCLEAN` the residuals are deviating above the $\pm 5\sigma$ threshold.

The measured residuals will certainly decrease while increasing the number of cleaning iterations within `tCLEAN`, but to the detriment of the computational cost. Employing the sample of 1000 ALMA cubes, we can compare the computational cost of the two algorithms when using 200 cleaning iterations for `tCLEAN`:

- **Blobs Finder made its predictions on the entire Test set in 23 seconds employing a single NVIDIA Tesla K20.**
- **tCLEAN took 4.3 minutes per cube utilising 8 Intel Xeon E5-2680 CPUs. Given the 400 CPUs at our disposal, we run it on 50 cubes at a time in parallel obtaining a total computational time of 1.5 hours**.

Employing BlobsFinder for the reconstruction task on the entire Test set results in a speed-up factor of 200 on our system with respect to `tCLEAN`. If we consider the possibility to accommodate `DeepFocus` within `tCLEAN`, if we account for the major cycle procedure execution time, this workout will provide an improvement in **speeding up the `CASA` procedures of at least a factor of 150**.

We are aware that the promising speed up in the procedures when employing `DeepFocus` within `tCLEAN` will not solve unfortunately the performance problem of the Wideband Sensitivity Upgrade [3] alone. However, it will make a huge contribution. In order to succeed for the ALMA2030 era, we want generalize our algorithm to perform a whole deconvolution process independently to `tCLEAN`. The pathway to implementing `DeepFocus` independently of `tCLEAN` and `CASA` is well-defined. However, securing the necessary funding is essential to hire the appropriate personnel for its successful execution.

Note on Speed Calculation and Sustainability: We plan to measure the execution speed of `DeepFocus` and `tCLEAN` in a setting suited for ALMA operations. We will also estimate and compare the associated costs.
So far, the speed comparison has been based only on prediction time. The total pipeline training time for `DeepFocus` was approximately 5 hours, while `tCLEAN` took around 1.5 hours (as `tCLEAN` does not require training time). However, this training time is a one-time investment when `DeepFocus` is used within the `CASA` framework. `DeepFocus` can be shared as a Python script along with a series of weight files (similar to pre-trained networks shared by Google, PyTorch, etc.), allowing the community to benefit from an active learning paradigm, where the model is periodically retrained on new and improved data to provide optimal performance.
`DeepFocus` operates on graphics processing units (GPUs) rather than central processing units (CPUs), potentially reducing long-term operational costs compared to using multiple CPUs. Infact, assuming we are buying hardware at the time of writing this document, a Intel Xeon or AMD EPY with 26 cores costs around 2000 euros. To reach the 400 cores employed in our computations, 16 of these must be clustered with a conservative estimated cost between 30,000 and 32,000 euros. A single NVIDIA A100 GPU costs between 10,000 and 12,000 euros. This is cheaper than setting up 400 CPU cores using high-performance CPUs, by roughly a factor of 3 which means that the computed speed-up factor, if price of hardware is accounted for, can be further augmented by this factor.

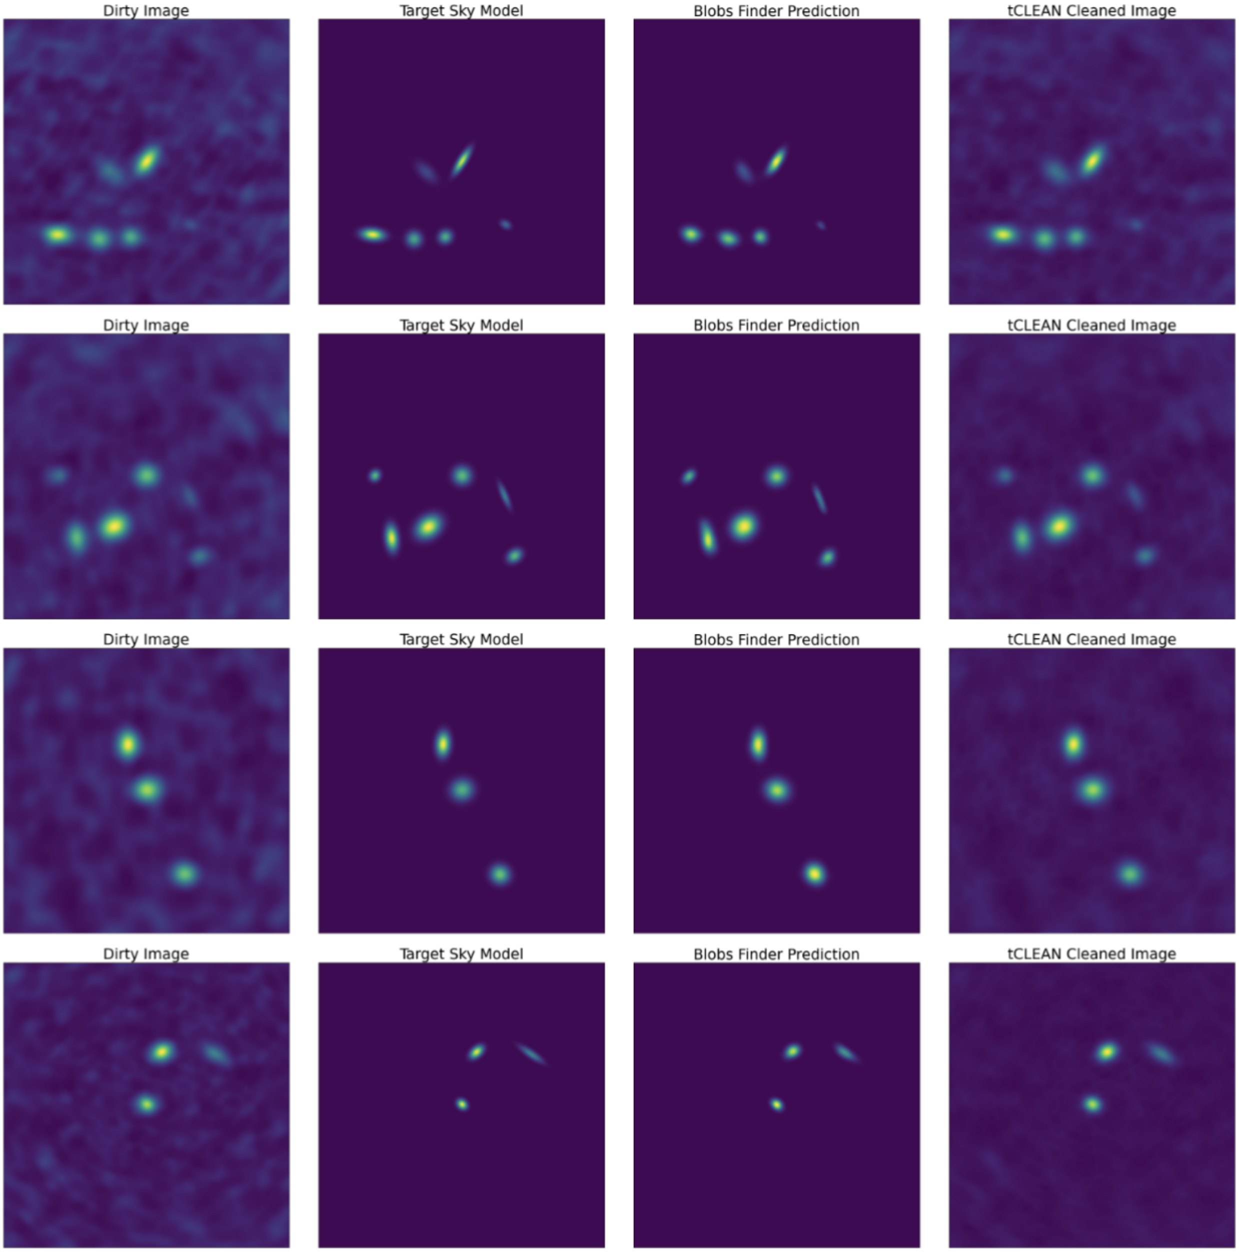

Figure 28: Comparison of `DeepFocus` and `tCLEAN` on a set of 1000 simulated ALMA cubes. In the images, the first, second, third and fourth columns shows a sample of Dirty Images, target Sky Models, Blobs Finder's reconstruction and `tCLEAN` reconstructions with niter = 200.

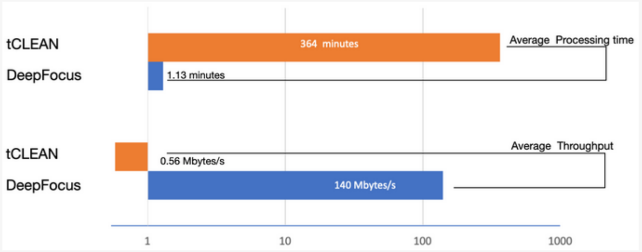

Figure 29: Comparison of processing time and computing throughput with `DeepFocus` and `tCLEAN` on 29×10^3 archived cube data from cycles 7, 8, and 9. This represents a rough estimate because at this stage of development, it is challenging to make a robust comparison between the two techniques.

5.5.2 Execution Speed Analysis on Archived ALMA Data Cubes

The European ALMA Regional Centre (EU ARC) cluster is regularly employed for processing the standard pipeline and generating official products delivered to principal investigators. Utilizing the EU ARC cluster, `DeepFocus` was executed on 29×10^3 archived ALMA data cubes collected during the three most recent observation cycles (Cycles 7–9; see Fig. 29).

For this experiment, `DeepFocus` was trained and tested primarily on simulated data, with limited use of real sources. At the time of its development, using only real data was not feasible due to the reliance on supervised learning methods and the ALMA archive's limitations in efficiently providing large volumes of data. Retrieving sufficient real data would require thousands of requests via the helpdesk and extensive processing with `CASA`, averaging 1–2 hours per cube, rendering the process impractical. Moreover, training directly on real data could introduce challenges due to uncertainties inherent in the data. Future integration with `ALMASim` aims to address these issues by enabling faster simulation generation and more efficient data retrieval from the ALMA archive. Additionally, no `DeepFocus` outputs on real data were recorded during this experiment. However, this will become feasible once `ALMASim` is fully developed and operational.

`DeepFocus`'s average processing time per cube was 1.13 minutes, with a compute throughput of 140 MB/s. For comparison, the same dataset was processed using the `tCLEAN` algorithm [12], [11], with parameters set to niter=1000, corresponding to 1000 cleaning iterations, and parallel computing enabled. The average compute throughput with `tCLEAN` was 0.56 MB/s, resulting in a performance rate approximately 250 times lower than that achieved by the `DeepFocus` ML algorithm.

The speed improvement offered by `DeepFocus` varied depending on the image size, with speedups ranging from **280-fold to an exceptional 5500-fold increase**. When using advanced algorithms like `DeepFocus`, the image deconvolution process can be completed within minutes, even for large data cubes. The use of GPUs, commonly employed in ML algorithms, provides significant benefits for synthesis image analysis. `DeepFocus` demonstrated both high image fidelity and impressive computational performance in reconstructing images from ALMA data cubes.

5.6 Bayesian Optimization and Meta–Learning for `DeepFocus` Model Selection

The `DeepFocus` deep learning pipeline [2] performs deconvolution, source detection, and characterization. Originally developed to detect faint compact objects Figs. 19, 30, the algorithm has been enhanced to detect extended emissions through the integration of a new ML model: The meta–learner Fig. 31. This meta-learner explores various architectures (e.g., CAE-VAE, U-Net, and ResNet) and aids in selecting the best-performing model for a given task and set of interferometric data. A Bayesian optimization algorithm is employed for model selection, supported by a taxonomy that incorporates multiple architectures, hyperparameters, and evaluation metrics specific to the problem, data, and desired performance criteria.

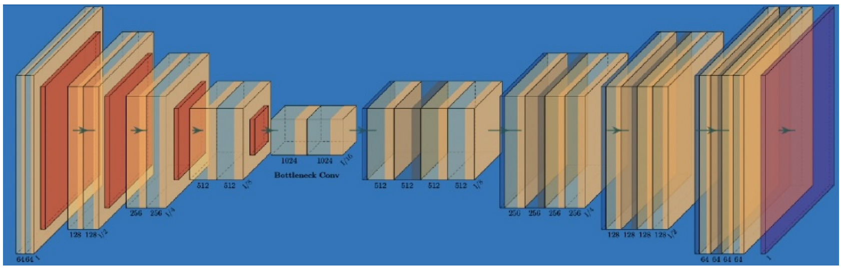

Figure 30: *An autoencoder is a type of neural network architecture that uses an encoder (right) to compress an input into a lower-dimensional representation (center), and a decoder (left) to reconstruct the original input from the compressed representation.*

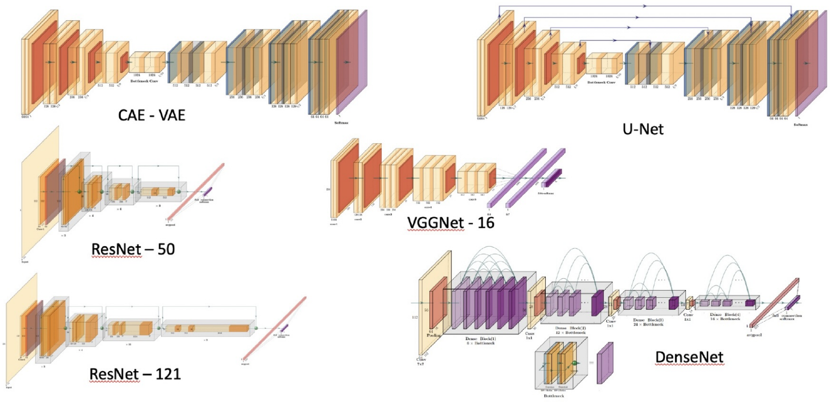

Figure 31: *The deconvolver `DeepFocus` is a meta–learner that searches for the most performing architecture given a task and a set of interferometric data.*

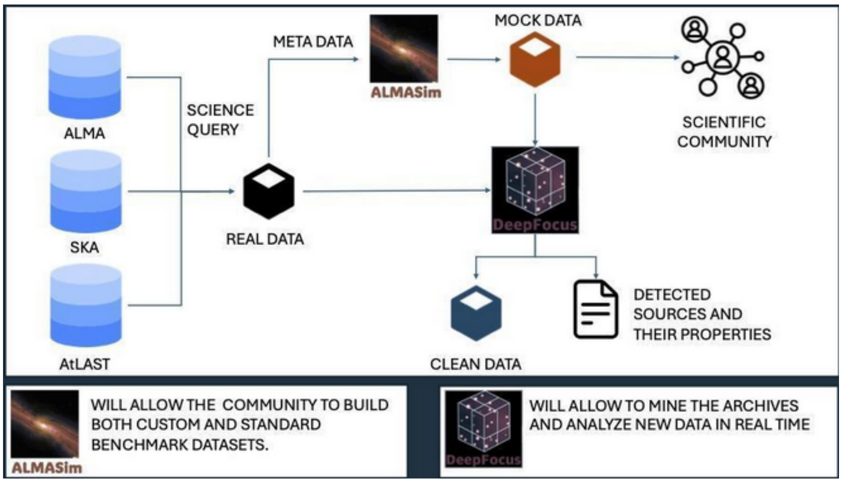

Figure 32: Schematic view of the described software components interacting with radio-interferometric Science Archives. `DeepFocus` *and* `ALMASim` *are flexible to be extended to other observatories.*

During optimization, multiple parameter realizations are tested in parallel on subsamples of the data to measure performance. Once the optimal architecture is identified, the pipeline handles tasks such as deconvolution, denoising, focusing, and classification.

The Bayesian optimization algorithm efficiently identifies the best set of parameters by reducing the number of expensive model evaluations through the use of surrogate models. The deep learning model (the objective function) is approximated by a Gaussian process (GP) surrogate, where $y = GP(\mu, \Sigma)$, with $\mu(x)$ and $\Sigma(x, x')$ representing the mean and covariance functions, respectively. A Matérn kernel is used for the covariance function to capture smoothness and correlations, defined as $\Sigma = K(x, x') = \sigma^2 \left(1 + \frac{(x-x')^2}{2\alpha l}\right)^{-\alpha}$ where α and l are smoothness and length scale parameters.

The surrogate model is trained on a limited number of initial objective function evaluations. An acquisition function, in this case, the expected improvement (EI), guides the selection of the next parameter set to evaluate. The EI function is defined as $EI(x) = \mathbb{E}\left[\max\left(f(x) - f(x'), 0\right)\right]$, where $f(x)$ and $f(x')$ represent the predicted mean of the objective function and the best value observed so far. This approach balances exploration, favoring high uncertainty $\sigma(x)$ and exploitation, preferring values where $\mu > f(x')$. As demonstrated in [86], surrogate models significantly reduce the number of costly objective function evaluations during Bayesian parameter optimization.

In summary, for a given dirty data cube, `DeepFocus` employs a taxonomy-based approach to identify the most suitable models for the specific dataset. The process begins with deconvolution and denoising to enhance the data quality. Following these steps, the algorithm performs focusing, regression, and classification to extract and characterize the relevant features from the data.

5.7 Summary and Outlook

`DeepFocus` offers efficient cube imaging, source extraction, and characterization for large datasets with minimal impact on data storage, thanks to its Meta-Learner that optimizes deep learning architectures. It has shown remarkable performance improvements through advancements in data loading, parallel processing, and GPU usage.

Although ALMA2030 [3] is still in early stages, `DeepFocus` is expected to have significant applications both within and outside ALMA operations. It could enable real-time processing and provide comprehensive imaging of all requested spectral windows without the need for traditional mitigation processes. Nevertheless Deep

Focus was only trained and tested on cubes containing at least a source, the classification part of its pipeline works on cube patches which may contain no source at all. This means that the application of Deep Focus on data cubes with no source inside should not generate false positives. To prove this, further experiments will be conducted leveraging ALMASim's capabilities of generating simulations without a source. Furthermore we are studying the possibility of using simulations with no source to study to characterize the noise properties of the simulations. `DeepFocus` can be integrated into the next generation of `CASA` as a deconvolution task or used as a standalone tool by any interested users or Principal Investigators. As shown in Fig. 32, `DeepFocus` can be advanced to learn from freshly archived data. For this purpose, `DeepFocus` will be trained as the archive gets updated. `ALMASim`, in conjunction with real data, serves as a crucial link for the scientific community, facilitating the development and evaluation of innovative ML techniques. It enables comprehensive quality assessment by allowing researchers to regenerate simulated data for comparative analysis. This capability ensures that results obtained with simulated data by one group can be replicated and compared with results from other techniques by different researchers. Both real and mock data are processed by `DeepFocus` to generate clean outputs and produce catalogs of detected sources.

Regarding the application of Deep Focus on complex line shapes, we haven't tested the capabilities of DeepGRU in reconstructing complex spectral profiles. The use case was never tested. That being said, DeepGRU was only one of the models explored in the taxonomy of models tested by DeepFocus. When Deep Focus will be applied to complex line profiles, a different model (more modern) could rise up to be the one suited to handle them. Regarding Deep Focus training and testing on simulated sources only. Employing real data for training was not feasible at the time of development of Deep Focus given that most algorithms at the time worked in a supervised fashion and given hat the ALMA archive does not offer a service to obtain the volume of needed data in reasonable time. Even if a user is capable of producing the thousands of data requests needed to gather enough trianing data, this data must be processed with `CASA` (to be employed for ML) which takes on our experience between 1 and 2 hours per cube. This would end up in an unreasonable time to obtain a meaningfully training set. Furthermore, the direct employment of real data for training could be a double-edged sword given the uncertainty on the data itself. Through ALMASim in the future we will be able to hopefully cover both the fast generation of simulations and the fast retrieval of real data from the archive.

The application on complex morphological profiles to address the detection of extended emissions has encountered several challenges, primarily due to the current limitations in ML techniques within our community. Large-scale applications of ML methods are constrained by several factors:

1. **Simulation Complexities**: Existing tools for generating simulations are inadequate for producing realistic observational signatures and often struggle to generalize to real-world scenarios. These tools are frequently proprietary, limiting access and collaboration among users. Additionally, the simulated data often come with incomplete metadata and lack a user-friendly interface, requiring significant time investment from experts, especially when thousands of simulations are needed for ML applications.

2. **Availability and Structure**: Simulators are typically not designed to support extensive and varied use cases, leading to a narrow focus that does not accommodate the broader needs of ML applications. This limitation exacerbates the difficulty in adapting simulations for diverse scenarios.

3. **Fallback Issues**: The absence of standardized benchmarks and the tendency for scientists to conduct independent efforts in both simulations and ML applications contribute to the fragmentation of research. This lack of common benchmarks impedes the comparison and validation of different techniques and results.

Consequently, our efforts have been concentrated on advancing `ALMASim`. Detailed information on the enhancements made to `ALMASim`, including the integration of a graphical user interface (GUI) and improvements to its accessibility for the scientific community, is provided in Section 6. `ALMASim` generates spectral profiles for various source classes, simulating luminosity functions for 180 species commonly observed in ALMA data, such as the CO (2–1) transition, enhancing the quality of our simulations. To refine the characterization of noise in images and facilitate effective learning by ML algorithms, the empirical noise study has played a pivotal role in our BRAIN initiative. The prototype version of `NOISEMPIRE` has been incorporated into `ALMASim`, with further details available in Section 7. Additionally, synergies among core team members have led to the integration of `NIFTy` into `ALMASim`. These advancements have been essential and have demanded a considerable amount of dedication and effort.

An archived ALMA QSO sample is prepared for input into the `DeepFocus` algorithm to search for faint serendipitous galaxies. While improving the detection of faint objects, efforts will focus on enhancing the algorithm's robustness. Previous work with datasets like SKA Data Challenge 2 and Westerbork Synthesis Radio Telescope data suggests that false detections, common with methods like Deep GRU and Blobs Finder, can be mitigated by post-detection classifiers. ResNets and Random Forests have already proven effective for this task.

Tests will verify the algorithm's ability to detect off-axis sources across channels, including extreme cases. Larger data cubes will be created and tested to address challenges posed by upcoming ALMA upgrades. The method will also be trained to detect continuum in continuum imaging and isolate noise in cubes and continuum images.

We want to refine the simulations employing physical models for the galaxy kinematics. Refinements on the simulation code and detection algorithm are also needed to detect other complex and reach environments as when observing the galactic center or the SZ effect for decrement features.

Because of the several and complex celestial radio sources detectable with ALMA, the algorithm shall be expanded to allow for transfer learning. For instance, domain adaptation will allow the algorithm's ability to learn from both artificial and real data. The basis of this development are important in operations in view of ALMA2030 [3].

6 ALMASim

`ALMASim` [87] is an open source simulator developed in Python and distributed through GitHub and PiP which aims to provide a reliable and user-friendly tool for generating realistic ALMA mock data based on real metadata from the ALMA TAP archive. Designed for both astronomers and computer scientists, `ALMASim` offers flexibility in its usage. Users can opt for a fully automated mode, requiring no prior knowledge of radio interferometry, or can customise simulations for specific ALMA sources and observing conditions. `ALMASim` is adaptable to various computational environments, working either on laptops and high-performance computing (HPC) clusters. It supports sequential execution for resource-constrained scenarios and leverages MPI programming for efficient parallel processing to generate large volumes of simulated data rapidly. The `ALMASim` Graphical User Interface, shown in Fig. 33, allows the user to setup and customize the simulation on a local machine or on a remote High Performance Cluster leveraging one of three possible schedulers: Dask, Slurm or MPI. The `ALMASim` simulation pipeline can be divided into five main phases:

1. **Simulation Setup**: In this phase, the user provides information about the desired simulations, including the number of simulations, the types of sources to simulate (sources skymodels), and the observing conditions.
2. **Metadata Retrieval**: Based on user preferences, observational metadata is retrieved from the ALMA TAP archive. This can be done by specifying a list of known targets or by using keywords available in the ALMA Archive such as scientific category, observing band, or field of view. For each target arising from the SQL query the following information are recorded: Band, PWV, RA, DEC, Int Time, Cont Sens, Angular Resolution, Bandwidth, Frequency of observation, Frequency Support (all SPWs) and the antenna_array
3. **Source Reasoning**: The retrieved metadata is then used to calculate relevant information for the simulations, such as the expected radio continuum brightness, the types of spectral lines expected in the given

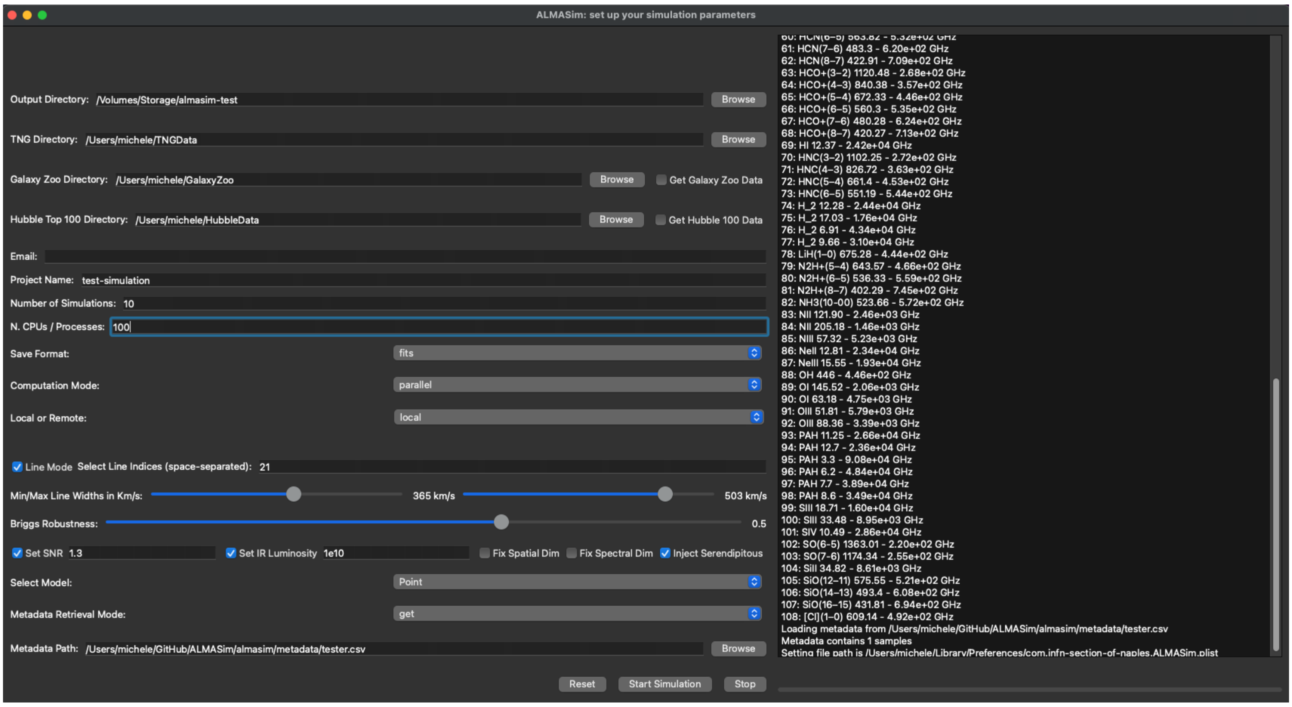

Output Directory: /Volumes/Storage/almasim-test

TNG Directory: /Users/michele/TNGData

Galaxy Zoo Directory: /Users/michele/GalaxyZoo ☐ Get Galaxy Zoo Data

Hubble Top 100 Directory: /Users/michele/HubbleData ☐ Get Hubble 100 Data

Email:

Project Name: test-simulation

Number of Simulations: 10

N. CPUs / Processes: 100

Save Format: fits

Computation Mode: parallel

Local or Remote: local

☑ Line Mode Select Line Indices (space-separated): 21

Min/Max Line Widths in Km/s: 365 km/s 503 km/s

Briggs Robustness: 0.5

☑ Set SNR 1.3 ☑ Set IR Luminosity 1e10 ☐ Fix Spatial Dim ☐ Fix Spectral Dim ☑ Inject Serendipitous

Select Model: Point

Metadata Retrieval Mode: get

Metadata Path: /Users/michele/GitHub/ALMASim/almasim/metadata/tester.csv

```
60: HCN(6–5) 563.82 - 5.32e+02 GHz
61: HCN(7–6) 483.3 - 6.20e+02 GHz
62: HCN(8–7) 422.91 - 7.09e+02 GHz
63: HCO+(3–2) 1120.48 - 2.68e+02 GHz
64: HCO+(4–3) 840.38 - 3.57e+02 GHz
65: HCO+(5–4) 672.33 - 4.46e+02 GHz
66: HCO+(6–5) 560.3 - 5.35e+02 GHz
67: HCO+(7–6) 480.28 - 6.24e+02 GHz
68: HCO+(8–7) 420.27 - 7.13e+02 GHz
69: HI 12.37 - 2.42e+04 GHz
70: HNC(3–2) 1102.25 - 2.72e+02 GHz
71: HNC(4–3) 826.72 - 3.63e+02 GHz
72: HNC(5–4) 661.4 - 4.53e+02 GHz
73: HNC(6–5) 551.19 - 5.44e+02 GHz
74: H_2 12.28 - 2.44e+04 GHz
75: H_2 17.03 - 1.76e+04 GHz
76: H_2 6.91 - 4.34e+04 GHz
77: H_2 9.66 - 3.10e+04 GHz
78: LiH(1–0) 675.28 - 4.44e+02 GHz
79: N2H+(5–4) 643.57 - 4.66e+02 GHz
80: N2H+(6–5) 536.33 - 5.59e+02 GHz
81: N2H+(8–7) 402.29 - 7.45e+02 GHz
82: NH3(10-00) 523.66 - 5.72e+02 GHz
83: NII 121.90 - 2.46e+03 GHz
84: NII 205.18 - 1.46e+03 GHz
85: NIII 57.32 - 5.23e+03 GHz
86: NeII 12.81 - 2.34e+04 GHz
87: NeIII 15.55 - 1.93e+04 GHz
88: OH 446 - 4.46e+02 GHz
89: OI 145.52 - 2.06e+03 GHz
90: OI 63.18 - 4.75e+03 GHz
91: OIII 51.81 - 5.79e+03 GHz
92: OIII 88.36 - 3.39e+03 GHz
93: PAH 11.25 - 2.66e+04 GHz
94: PAH 12.7 - 2.36e+04 GHz
95: PAH 3.3 - 9.08e+04 GHz
96: PAH 6.2 - 4.84e+04 GHz
97: PAH 7.7 - 3.89e+04 GHz
98: PAH 8.6 - 3.49e+04 GHz
99: SIII 18.71 - 1.60e+04 GHz
100: SIII 33.48 - 8.95e+03 GHz
101: SIV 10.49 - 2.86e+04 GHz
102: SO(6-5) 1363.01 - 2.20e+02 GHz
103: SO(7-6) 1174.34 - 2.55e+02 GHz
104: SiII 34.82 - 8.61e+03 GHz
105: SiO(12–11) 575.55 - 5.21e+02 GHz
106: SiO(14–13) 493.4 - 6.08e+02 GHz
107: SiO(16–15) 431.81 - 6.94e+02 GHz
108: [CI](1–0) 609.14 - 4.92e+02 GHz
Loading metadata from /Users/michele/GitHub/ALMASim/almasim/metadata/tester.csv
Metadata contains 1 samples
Setting file path is /Users/michele/Library/Preferences/com.infn-section-of-naples.ALMASim.plist
```

Figure 33: The ALMASim Graphical User Interface.

band, their fluxes, velocities, redshifts, and signal-to-noise ratios (SNR). This allows for the simulation of skymodels which have similar characteristics to those of the corresponding real observations.

4. **Sky Model Simulations**: Using the source information, the incoming signals are simulated in a noiseless environment, creating ”Sky Models.” ALMASim currently supports Point, Gaussian, Extended, and Diffuse sky models. Extended sky models are generated by incorporating Illustris TNG [21] [88], Galaxy Zoo [89] or Hubble [90] images into the data cubes.
5. **ALMA Simulation**: The collected metadata and the generated Sky Models are fed in input into a radio-interferometric simulator to produce a realistic observation of the Sky Model. The observational characteristics such as the antenna configuration, the beam size, the integration time, the field of view, the continuum sensitivity and so forth are derived from the corresponding real observation metadata. The resulting observed cube (Dirty Cube), the dirty visibilities, and model visibilities (Fourier inversions of the Dirty and Sky Model cubes, respectively) are then, if possible, compared with the real counterpart to check for errors and saved on disk.

The relevant parameters for the simulation setup are the followings:

1. **Output Directory**: a text field and a browse button to specify and select the directory where the simulation results will be saved.
2. **TNG Directory**: a text field and a browse button to specify and select the directory where the Illustris TNG [21] data will be downloaded (if needed).
3. **Galaxy Zoo Directory**: A text field and a browse button to specify and select the directory where Galaxy Zoo data will be downloaded (if needed).

4. **Hubble Top 100 Directory**: A text field and a browse button to specify and select the directory where Hubble data will be downloaded (if needed).

5. **Email**: a text field for the user email address. The email address is used to notify the user of the completion of the simulation if they are run on a remote machine.

6. **Project Name**: a text field to enter the name of the simulation project. This will be a sub-directory created within the Output Directory where all the data products will be stored. If this field is left empty, data will be stored in the Output Directory.

7. **Number of Simulations**: a text field to enter the number of simulations to run.

8. **Total Number of CPUs**: a text field to enter the total number of CPU cores involved in the simulations.

9. **Save Format**: A drop-down menu to select the format in which to save the simulation results. The supported formats are *.fits*, *.npz* and *.h5*.

10. **Computation Mode**: a drop-down menu to select the computation mode. Allowed modes are *sequential* and *parallel*. In sequential mode, simulations are run sequentially using the total number of selected CPUs, while in parallel mode the Dask [91] is used to parallelise over multiple CPUs.

11. **Remote Mode**: a drop-down menu to select if the simulations have to be run on the local machine or remotely. In the latter case, the following additional fields are presented to the user:
    (a) **Remote Computational Mode**: a drop-down menu showing three possible configurations: *MPI*, *Slurm* and *PBS*. If *MPI* is selected, then a Dask cluster is created on the remote machine and used to perform the simulations in parallel over the selected amount of CPUs. If *Slurm* is selected, a Dask SlurmCluster is created on the remote machine and used to perform the simulations. A Dask SlurmCluster is a tool that bridges the Dask parallel computing library with Slurm [92], a workload manager commonly used on high-performance computing (HPC) cluster, and allows to directly leverage the resources managed by Slurm. If *PBS* is selected, Dask interfaces itself with the Portable Batch System (PBS) workload manager on the remote machine;
    (b) **Remote Host**: the host-name or IP address of the remote machine;
    (c) **MPI/Slurm/PBS Config**: a .json file containing information about the cluster, such as FQDN of the scheduler, ports and available resources;
    (d) **Username**: the user name on the remote machine;
    (e) **SSH Key**: the SSH key to connect to the remote machine with password-less ssh;
    (f) **Key Password**: the password for decrypting the SSH Key (if needed);
    (g) **Set Work Directory**: a checkbox to control whether the simulations and all the necessary ancillary files should be stored in the user remote *home* folder or in a different working directory. If the checkbox is ticked, the user is asked for the full path of the working directory on the remote machine.

12. **Line Mode**: a checkbox to enable or disable line mode. If line mode is enabled, a list of 108 lines is displayed, and the user is asked to select which lines to simulate. If the line mode is disabled, the user is asked for the redshift distribution of the simulated observations, and the number of lines that should fall in the observation spectral window (e.g. the frequency range observed by ALMA).

13. **Min/Max Line Widths in km/s**: two sliders controlling the minimum and maximum widths in km/s of the simulated emission lines.

14. **Brigg's Robustness Parameter**: a slider which allows to set the Brigg's robustness parameter for UV coordinates weighting.

15. **Set SNR**: a checkbox to control the Signal to Noise Ratio (SNR) of the simulated observations. If set, the user is asked to insert the SNR distribution of the simulated sources. If not set, the SNR is set randomly, but always higher than the expected ALMA sensitivity (e.g. all simulated sources should be detectable).
16. **Set IR Luminosity**: a checkbox to control the infrared luminosity normalisation. If set, the user is asked to insert the infrared luminosity to which the theoretical Spectral Energy Distribution (SED) is normalised to. If not set, the SNR is scaled to the infrared luminosity such that the minimum observed continuum is higher than the minimum continuum observable with the chosen ALMA configuration.
17. **Fix Spatial and Spectral Dims**: the two check-boxes control the size in pixels of the output data products. If set, the user is asked to insert the desired number of pixels and channels. If not set, the number of pixels are directly computed from the band and field of view sampled from the metadata.
18. **Inject Serendipitous**: a checkbox to control the injection of serendipitous sources. If set, a random number between 1 and 5 Gaussian sources are injected alongside the selected source model within the sky model cube. Each serendipitous source is simulated with the same line configuration as the central source, but with a lower SNR.
19. **Select Skymodel**: a drop-down menu to select the model to use. Seven source models are available, *Point*, *Gaussian*, *Extended*, *Diffuse*, *Galaxy Zoo*, *Hubble* and *Molecular*. The *Point* model is built by injecting a dot at a specific position (x, y, z) in the cube. The *Gaussian* model by injecting a 2D Gaussian Kernel at a specific position (x, y, z, FWHM_x, FWHM_y, pa) where FWHM_x, FWHM_y are the Full Width Half Maximums of the Gaussian Kernel in the *x* and *y* axes, and *pa* is the projection angle with respect to the line of sight. The *Extended* model is built by searching TNG snapshots (dictated by the desired redshift) for a compatible gas halo, and injecting it within the cube at a specific position (x, y, z). The *Diffuse* model is generated by injecting a 2D correlation random field (CRF) in each 2D plane of the cube. The CRF is created using a power spectrum based approach in which, first, a random correlation field is generated using [93], and then an exponential function is applied to it in order to ensure positivity and non-Gaussianity in the field. The *Galaxy Zoo* model is built by randomly sampling a Galaxy Zoo image and injecting it in all 2D planes of the cube. The resulting cubes are normalised, channel by channel, to the expected continuum and then convolved with 1D Gaussian Kernels (L, z, FWHM_z) to obtain the line emissions, where L is the line luminosity, and FWHM_z is the width of the line in number of channels. The *Hubble* model is built by randomly sampling an Hubble Top 100 image and injecting it in all 2D planes of the cube. The resulting cubes are normalised, channel by channel, to the expected continuum and then convolved with 1D Gaussian Kernels (L, z, FWHM_z) to obtain the line emissions, where L is the line luminosity, and FWHM_z is the width of the line in number of channels. The *Molecular Cloud* model is built by generating a 2D power-law image with a specified index and random phases and injecting it in all 2D planes of the cube. The resulting cubes are normalised, channel by channel, to the expected continuum and then convolved with 1D Gaussian Kernels (L, z, FWHM_z) to obtain the line emissions, where L is the line luminosity, and FWHM_z is the width of the line in number of channels.
20. **Metadata Retrieval Mode**: a drop-down menu to select the metadata retrieval mode. Two modes are available, *query* and *get*. If the *query* mode is selected, the ALMA Archive TAP Interfaces (EU, NA, EA) are queried for metadata on the basis of a list of targets or through a Scientific query (see Query Type below). If *get* is selected, the user is asked to select the file where metadata are stored on disk.
21. **Query Type**: a drop-down menu to select the scientific query mode. If *science* is selected a query is performed on the ALMA Archive on the basis of the following user choices: *scientific category*, *science keyword*, *band*, *frequency of observation*, *integration time*. If *target* is selected, the user is asked to point to a file containing the list of target names and ALMA archive identifiers (*group ID* and *member ID*), which are used to query the ALMA archive for relevant metadata.

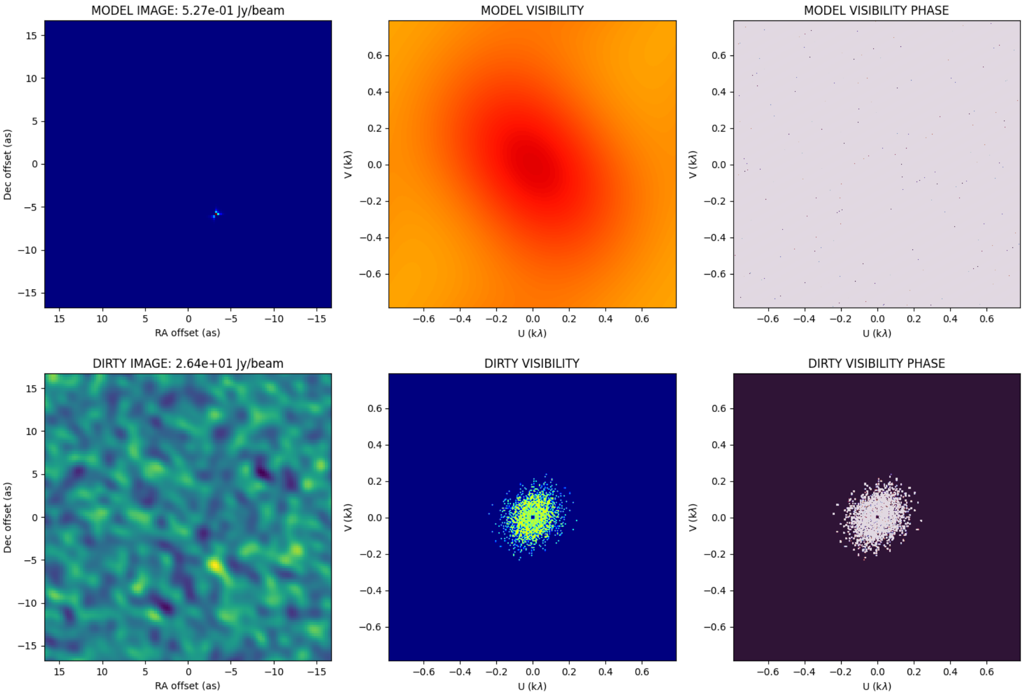

Figure 34: *Simulation of a Point Source with ALMASim.*

22. **Save/Load Metadata**: a text field and a browse button to specify and select the file where the queried metadata will be saved to (query mode) or loaded from (get mode).

Fig. 34, 35, 36, 37, 38, 39 show an example for each sky model that can be simulated with ALMASim. Efforts has still to be made in order to expand the simulator to employ radio catalogues as Radio Galaxy Zoo [94], VLA surveys (VLASS [95], FIRST [96]), Lofar surveys (LoTTS [97], MMSS [98]). This development is out of the scope of this study, but we are eager to address this topic in the future.

6.1 Simulating QSOs from the ALMA Archive

To generate realistic Quasi Stellar Objects (QSOs) luminosities and spectral profiles, ALMASim employes a Spectral Energy Distributions (SED), to generate the continuum, and a database of 110 line ratios to compute the line fluxes. The SED are produced by [99] through a "hybrid" model that combines a physical, forward model for proto-spheroidal galaxies with a phenomenological model for late-type galaxies. The authors use multi-wavelength data (from mid-infrared to millimeter waves) and introduce an evolving luminosity function for both AGNs and their host galaxies. The SED are normalized so that the IR luminosity, i.e. the integrated luminosity between 8 and 1000 μm, is 1 $L_{\odot}$. The two SEDs, normalized to $10^{10} L_{\odot}$ and shifted to redshift 0.05 are shown in Fig. 40. The database of line ratios is produced by Spilker et al. 2014 [100] which derived the average composite spectrum of a Dusty Star Forming Galaxy (DSFG) by combining the spectra of 22 DSFG observed by ALMA through a stacking method. The sources are shifted to the same redshift, $z = 3$,

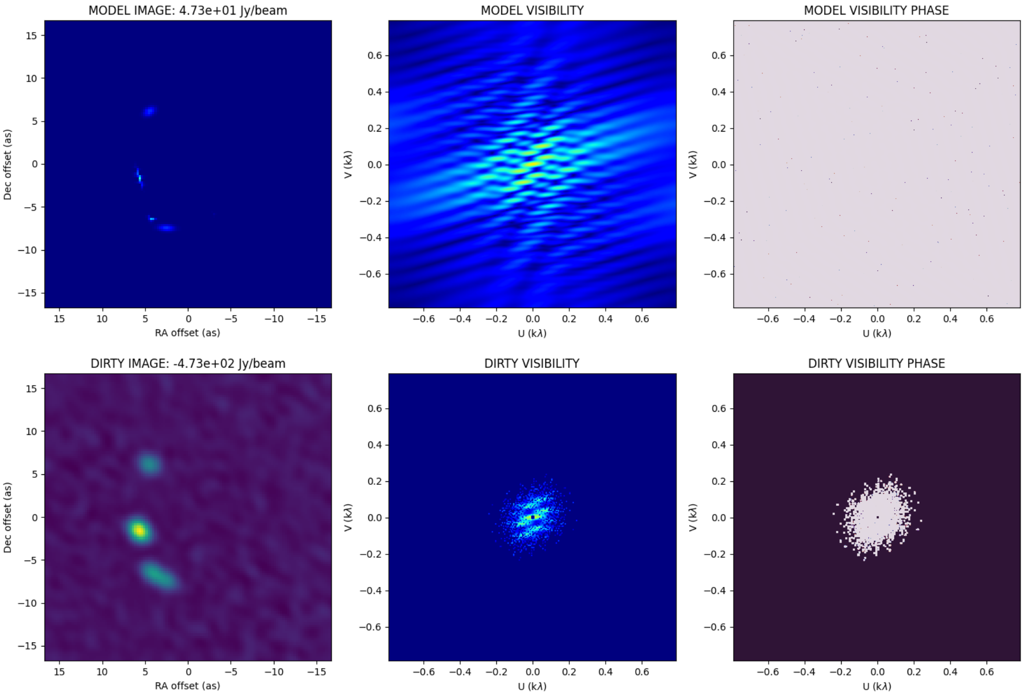

Figure 35: *Simulation of a Gaussian Source with multiple serendipitous sources with ALMASim.*

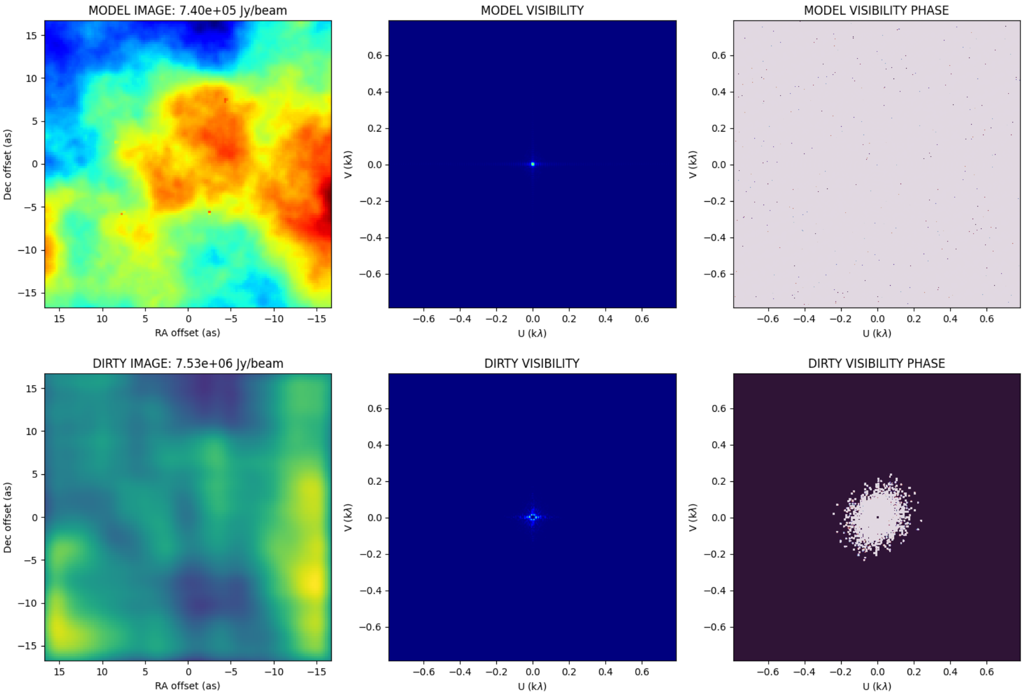

Figure 36: *Simulation of a Diffuse Field with ALMASim.*

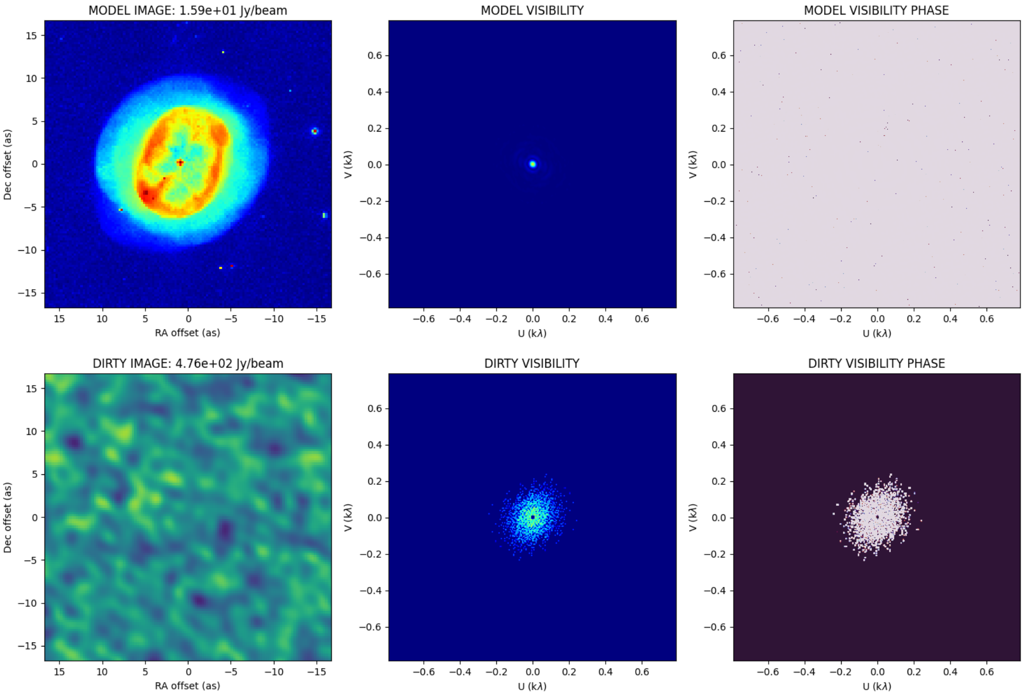

Figure 37: Simulation of a Hubble Top 100 Source with ALMASim.

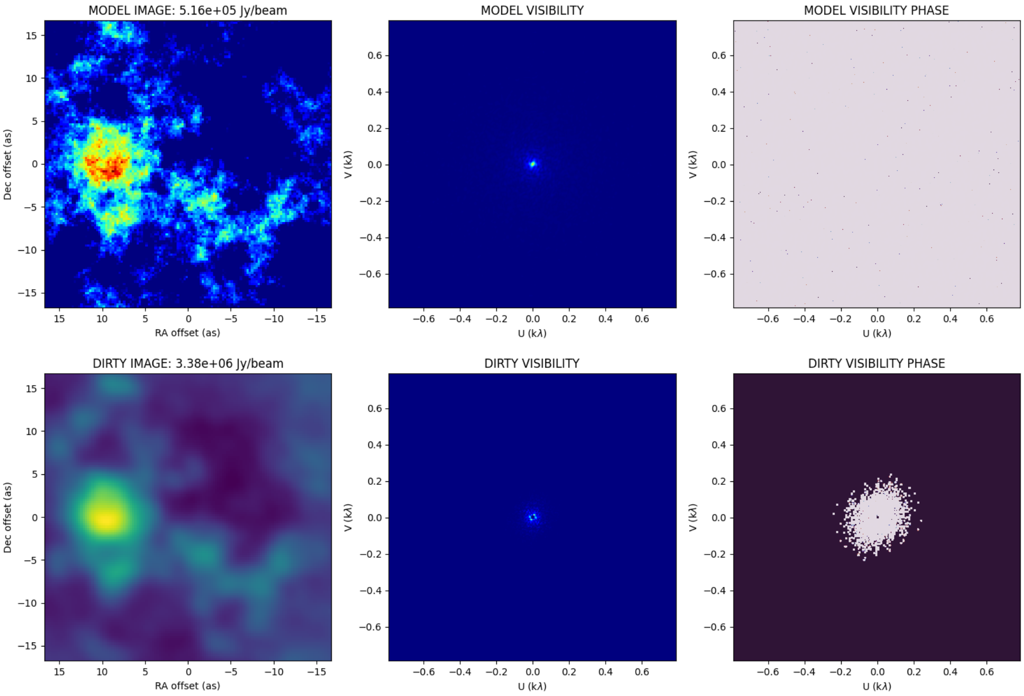

Figure 38: *Simulation of a Molecular Cloud with ALMASim.*

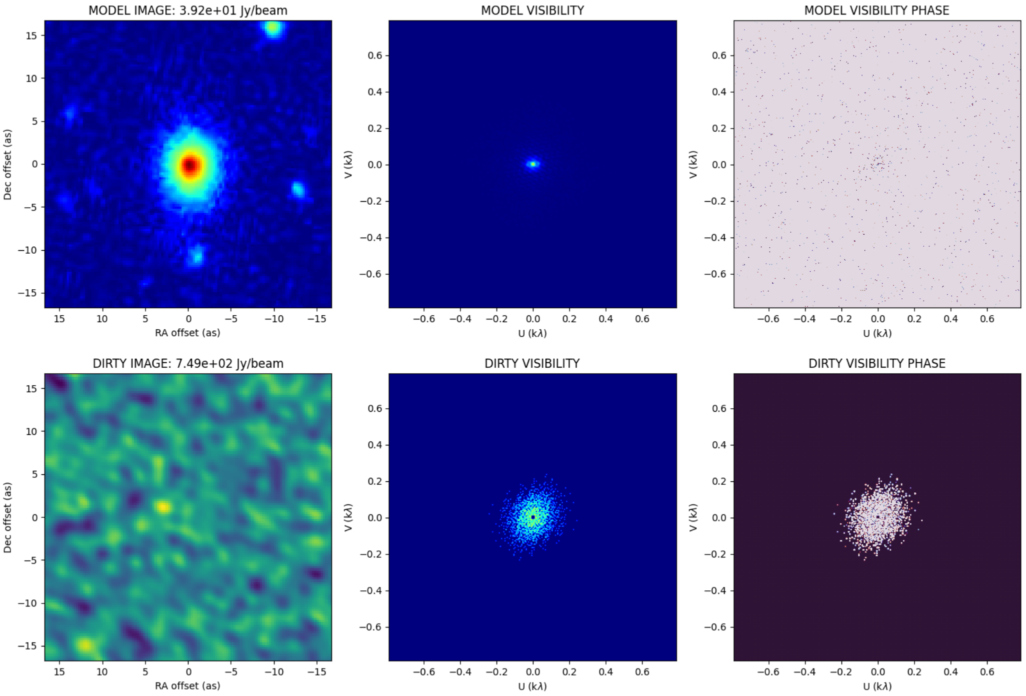

Figure 39: *Simulation of a Galaxy Zoo Galaxy with ALMASim.*

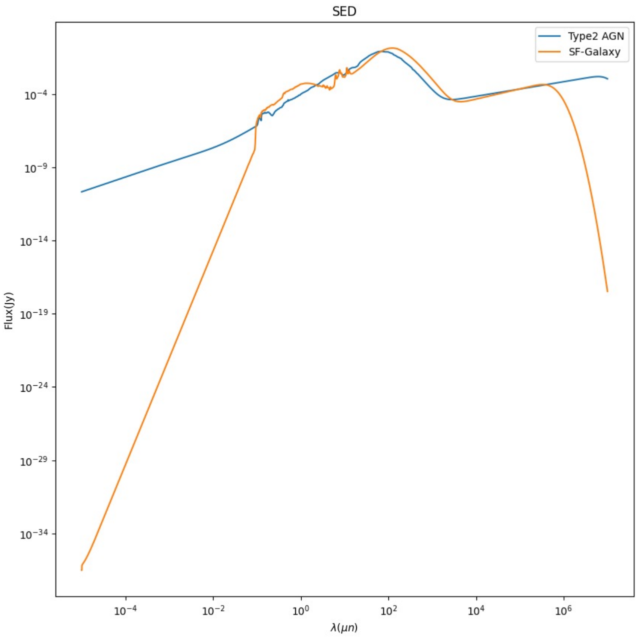

Figure 40: The two Spectral Energy Distribution employed by ALMASim to simulate source continuum

and rescaled to the same apparent infrared luminosity of 10^{13} solar luminosities. By analyzing the composite spectrum, several molecular lines were detected, such as 12CO, 13CO, HCN, HNC, HCO+, CN, even when not all transitions were individually visible in the individual spectra. Knowing the SED luminosity values of the AGN, and in particular their infrared luminosity, it is possible to trace the luminosity emitted by the different spectral lines through the relationship between the latter and the former via the parameter c, defined as the ratio between the line luminosity and a given infrared luminosity. Assuming that a line and an IR Luminosity are chosen, the SED is shifted in redshift such that the line falls within the SED, and the line fluxes are computed as follows:

$$log L_{line} = log L_{IR} + c \tag{5}$$

where:

L_{line} is the line luminosity integrated over the frequency step (or line width) in unit of Solar luminosity $L_\odot$,

L_{IR} is the infrared continuum in $L_\odot$ scale.

c the ratio between the line luminosity and a given infrared luminosity.

Since the given SEDs are not provided in Jy, but in $erg/s/Hz$, the relation used to convert them is :

$$\mathcal{L}[Jy] = \frac{L[erg/s/Hz] \times 10^{23}}{4\pi D_L^2[cm]} \tag{6}$$

The SEDs are re-normalized to the chosen infrared luminosity and line fluxes are computed as follows:

$$L_{line} = \frac{10^{\log(L_{IR})+c}}{\nu_{step}} \tag{7}$$

where ν_{step} are the line widths in Hz.

$$\nu_{step} = \frac{\Delta v}{c}\nu_{rest} \tag{8}$$

To assess the correctness of our method, we simulated the CO(1-0) transition line for a galaxy with a Infrared Luminosity of $L_{IR} = 5.0 \times 10^{13}$, and compared the results with those obtained empirically by [101]. Kirkpatrick et al. 2019 analysed the CO line emissions for a heterogeneous sample of galaxies distributed over a redshift range of 0.01 to 4. For the sake of comparison, we fixed a line width of $300km/s$ and changed the redshift between 0 and 5 and converted the luminosities computed by Kirkpatrick et al. in Jy through the following equation [102]:

$$L'_{CO} = \frac{c^2}{2k_B} S_{line} \Delta v \frac{D_L^2}{\nu_{obs}^2 (1+z)^3} \tag{9}$$

c is the velocity of the light

k_B is the Boltzmann constant

$S_{line}\Delta v$ is the measured flux of the line in $Jy\ km\ s^{-1}$

D_L is the luminosity distance in Mpc

ν_{obs} is the observed frequency.

that can be reorganized as

$$L'_{line} = 3.25 \times 10^7 \times S_{line} \Delta v \frac{D_L^2}{(1+z^3)\nu_{obs}^2} \tag{10}$$

Fig. 41 shows the line luminosities obtained from ALMASim employing the SED and the line ratios (blue) and those measured by Kirkpatrick et al. (red). The plot shows that the values produced by ALMASim are compatible with those measured by Kirkpatrick. A sample of 31 QSOs observed by ALMA in band 3, 4, 6, 7 and 8 was chosen as a test set to compare in detail `ALMASim` simulations and real observations. To that end, the QSOs Calibrated Visibilities were downloaded from the ALMA Archive and processed with `CASA tCLEAN` with 0 iterations to produce dirty cubes. `ALMASim` cannot perfectly reproduce a single source, given the uncertainty in the estimation of many source properties and observational parameters and conditions, but aims at producing a distribution of sources which should include the real ones. Figs. 42 and 43 show the integrated real dirty cubes (left), the simulated counterparts (right) and the spectral profiles of two QSOs among the 31. The integrated images were obtained by integrating the cube over the frequency axis, while the spectra by integrating each channel of the cubes over the spatial dimensions and by running a rolling mean function over the obtained spectra. Table 5 shows a comparison of the measured fluxes, signal to noise ratios and noise rms. To generate a dataset statistically similar to the real QSO sample, we generated 10 mock simulation for each cube obtaining a total of 310 mock data cubes. Fig. 44, Fig. 45, Fig. 46, Fig. 47, Fig. 48, Fig. 49, show, respectively, the distributions of source flux density, total cube flux, signal-to-noise ratio (SNR), root mean square (RMS) noise, beam size, and cell size measured respectively on the real cubes (orange) and mock ones (blue). Table 6 shows the measured distribution means and standard deviations and the p-values of the Kolmogorov-Smirnov (KS) tests employed to asses whether the distributions are statistically similar. A high KS value indicates a larger difference between the distributions, while a low value suggests that the distributions are quite similar.

The p-value indicates the probability of observing a difference as large as or larger than the one measured, assuming the two distributions are the same (null hypothesis). A high p-value (greater than 0.05) means there is insufficient evidence to reject the null hypothesis, implying that the two distributions are statistically similar.

6.2 Measuring Simulation Times and Parallelization Capabilities

To assess the parallelisation potential of `ALMASim`, we compared its computational times for simulating 1000 ALMA observations. We ran `ALMASim` in parallel mode on the IBISCO HPC (using 224 cores across 56 jobs, each with 4 cores and 16GB RAM) and in sequential mode on a laptop (with 4 cores and 16GB RAM). The cube

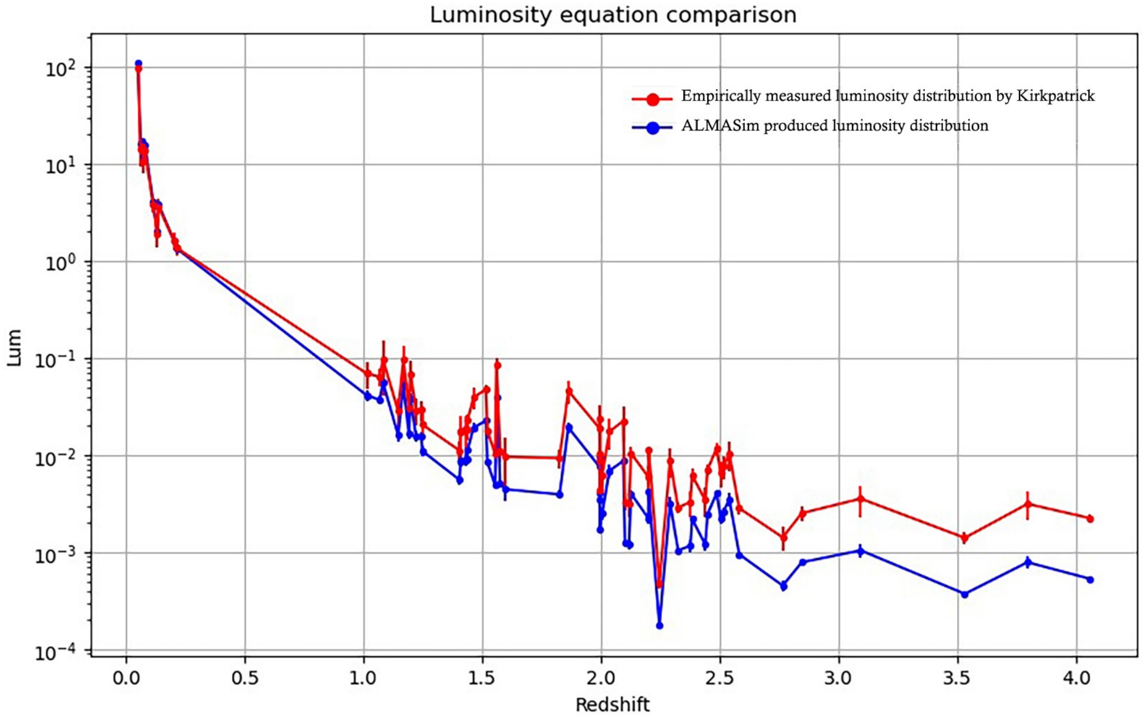

Figure 41: Comparison of the luminosity distribution produced by ALMASim employing SED and Line ratios (blue) and that empirically derived by Kirkpatrick et al. 2019 [101]

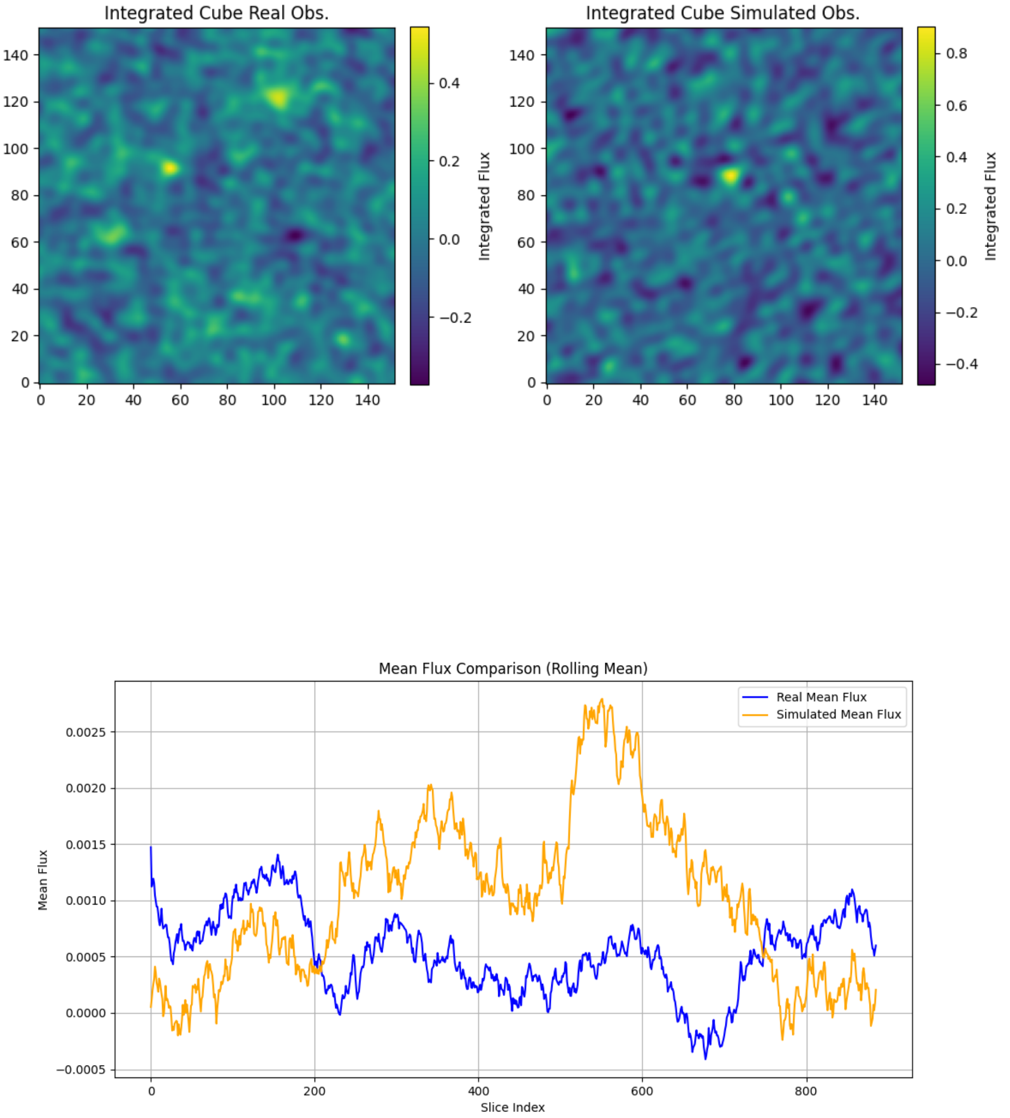

Figure 42: Top left: integrated image of the real observation of the QSO J0842. Top right: simulation made with ALMASim. Bottom: compared spectra of the two cubes.

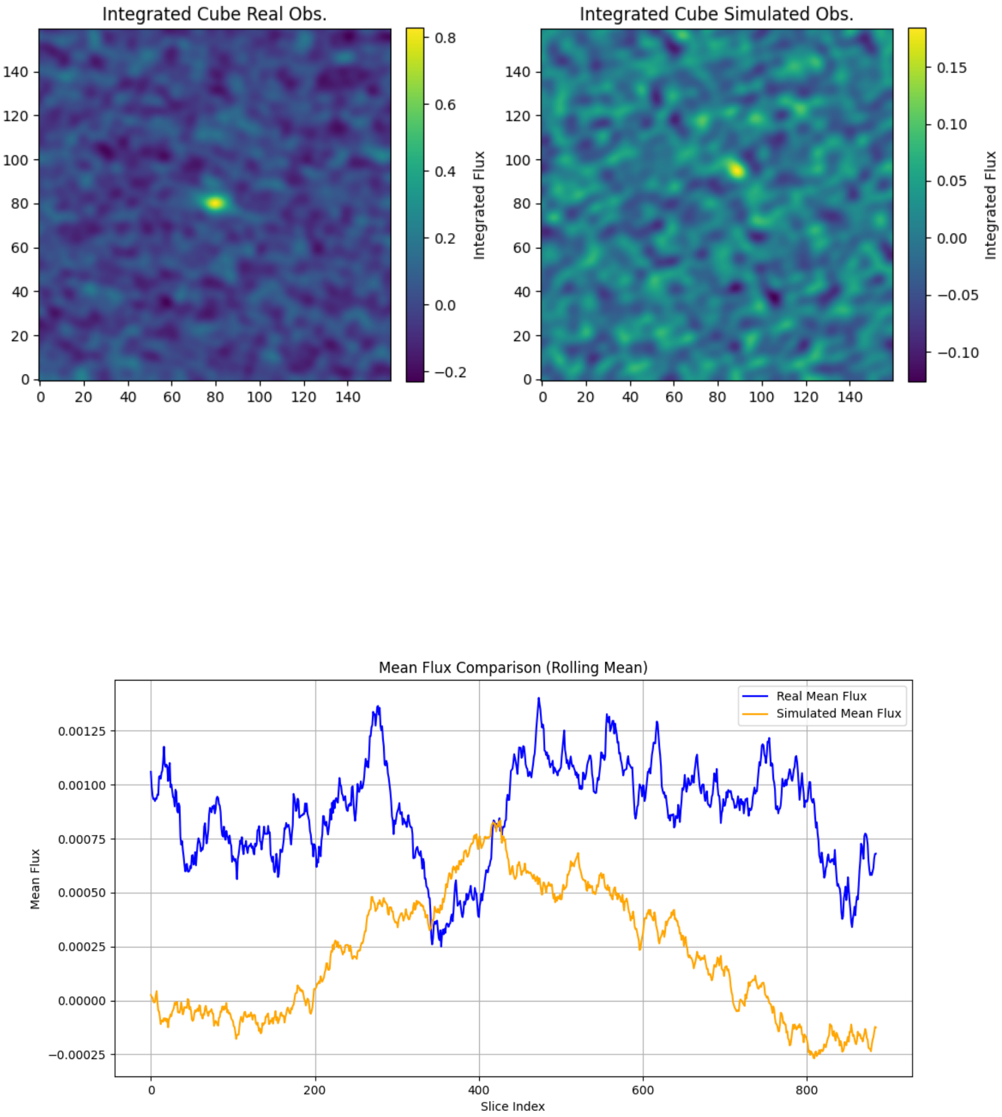

Figure 43: Top left: integrated image of the real observation of the QSO J1316+03564. Top right: simulation made with ALMASim. Bottom: compared spectra of the two cubes.

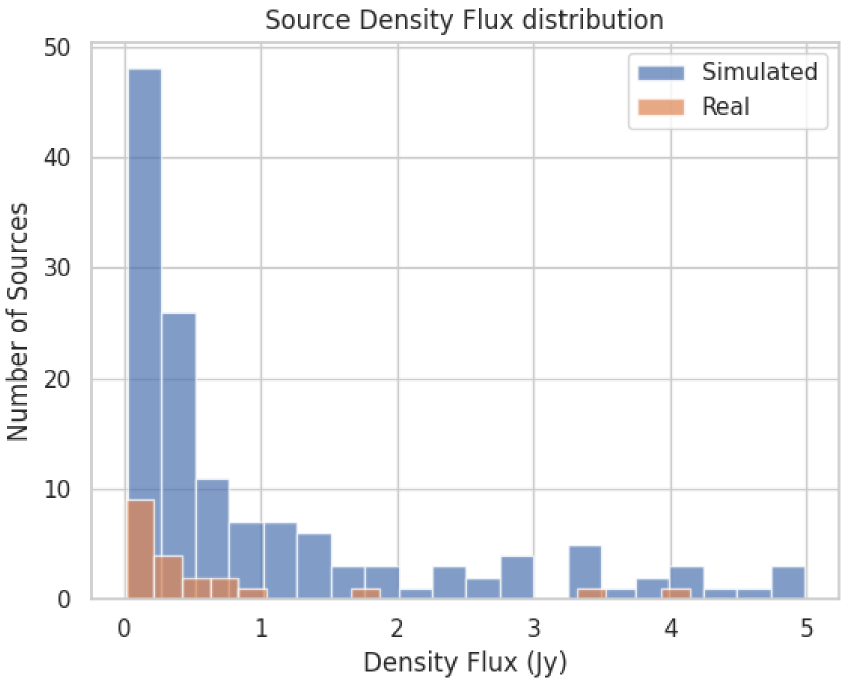

Figure 44: Comparison of the distributions of real (orange) and simulated (blue) source fluxes. Fluxes are measured by integrating around the source position in the x,y, freq space. The integration is carried within the fulls spectral extension of the cube and within the beam size in the spatial dimensions.

Property	J0842 Sim	J0842 Real	J1316 Sim	J1316 Real
Total Flux (Jy/beam)	6.92	16.37	8.06	-10.96
Source Flux (Jy/beam)	0.85	0.51	0.17	0.78
Noise RMS	0.0033	0.0020	0.0009	0.0022
SNR	6.82	8.09	6.42	7.54

Table 5: Comparison of ALMASim simulation and real cubes for J0842 and J1316+03564.

	Simulated Data	Real Data	KS	p-value
Source flux	1.06 ± 1.30	0.72 ± 1.12	0.19	0.45
Total flux	-36.32 ± 237.05	-51.93 ± 83.93	0.27	0.12
SNR	5.86 ± 2.34	5.86 ± 2.85	0.13	0.89
RMS Noise	0.12 ± 0.14	0.071 ± 0.097	0.34	0.023

Table 6: Mean, standard deviation, KS statistic and p-value measured on key parameters of the simulated and real QSOs. KS statistics and p-values indicates a high similitude between simulated and real data.

dimensions were set to $256 \times 256 \times 256$ pixels, and the source model was a single point source with an SNR of 1.5 and an infrared luminosity of $10^9 L_\odot$ solar luminosities. The metadata query was performed locally in both cases and excluded from the time analysis. The sequential run took 42.6 hours, averaging 2.56 ± 0.14 minutes per cube. The parallel run on IBISCO HPC took 47.32 minutes, averaging 2.83 ± 0.3 seconds per cube, while the average computation time per job on the HPC (2.65 ± 0.12 minutes) was higher than on the laptop. This is attributed to the overhead of exchanging metadata and messages via SSH between the user's local machine, the HPC UI, and the Slurm scheduler. Overall, we observed a speed-up factor of approximately 54. However, considering the 56-fold increase in resources used in the parallel run, **the computational efficiency on the HPC is estimated to be around** 96%.

6.3 ALMASim Final Remarks

With adequate development, `ALMASim` has the potential to serve as a versatile tool for various applications related to ALMA operations and research. Right now it can query the ALMA Archive, it can simulate data based on it, and we have a separate suit that employs code from `ALMASim` to feed `CASA` the necessary meta-parameters (beam size, cube dimension, antenna array) to transform the provided correlated visibilities in Deep Learning ready data cubes.

We want to expand the capabilities of the simulator to other facilities, to tackle total power and multi-wavelength astronomy. In case of ALMA total power data and new generation facilities, as AtLAST [103], our team is planning to combine `ALMASim` with the software `maria` [104] in collaboration with the AtLAST team with the aim of strengthening `ALMASim` to simulate single dish observations and to improve the atmospheric simulations for ALMA simulated datasets. This development is out of the scope of this study. We foresee that these implementations will allow `DeepFocus` to become capable to detect emissions from single dish observations, paving the way to address ALMA group imaging.

7 The empirical noise study

Noise significantly hinders astronomical image interpretation. Current imaging tools are complex and time-consuming, a problem that will be exacerbated by the upcoming data deluge from facilities like ALMA after the WSU or SKA. ML techniques like `DeepFocus` promise to offer a solution, but they often require to be trained on simulations as realistic as possible. On the other hand, realistic simulations are also necessary every time

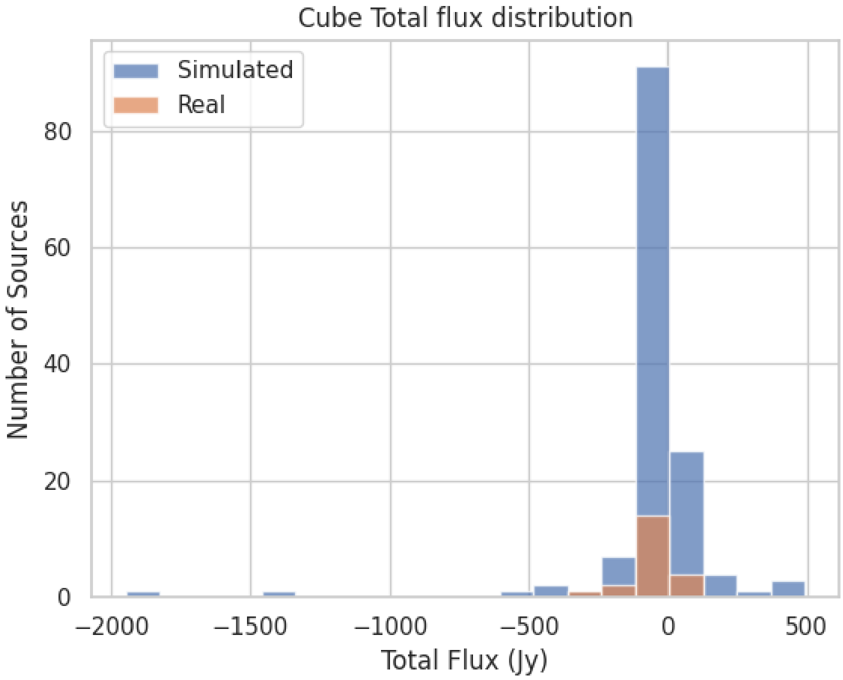

Figure 45: Comparison of the distributions of real (orange) and simulated (blue) total fluxes injected in the cubes.

the reliability of faint detection needs to be assessed (through completeness and contamination studies), for example when computing number counts [105].

Current simulation algorithms (such as the `CASA` internal tasks “simalma” and “simobserve”) struggle to accurately replicate the complex noise patterns found in interferometric astronomical images. These algorithms rely on theoretical models that can’t fully capture the complex characteristics of real-world noise. Noise in interferometric images tends to be spatially correlated and displays non-stochastic patterns and flux distributions, meaning the pixel fluxes don’t follow a Gaussian distribution. Even complex theoretical noise models considering random (statistical data distribution) noise, instrumental noise (including electronics), cosmic noise and systematic errors introduced by the calibration process always miss some unpredictable features and patterns observable in real images. To create more realistic simulations, it’s essential to analyze real images identifying and replicating the specific noise patterns observed in them.

We developed a new method (`NOISEMPIRE`, Baronchelli et al. 2024, in preparation) to create more realistic simulations of astronomical images. Unlike traditional methods, `NOISEMPIRE` analyzes real images to capture the specific noise patterns and replicate them in simulations. This empirical approach is particularly useful for interferometric data (like from ALMA) but has the potential to work with other instruments and wavelengths. The current version of `NOISEMPIRE` is a Python-based prototype freely available online (version 1.0.2: available on GitHub at this link). By combining different techniques, `NOISEMPIRE` is currently capable of:

- **Isolating sources**: The software internally runs source-extractor (`SExtractor` [106]), automatically determining the most effective combination of parameters.
- **Measuring the background at different scales**: Depending on the PSF of the input image, background is computed at different scales.
- **Measuring the RMS locally**: The local simulated RMS mimics the variability observed in the input image.

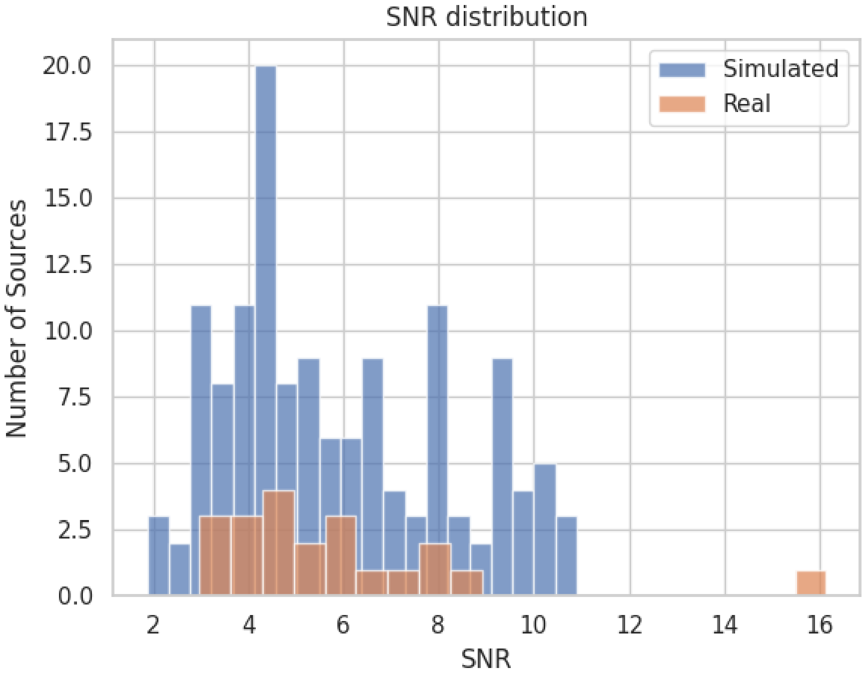

Figure 46: Comparison of the distributions of real (orange) and simulated (blue) Signal to Noise Ratios (SNR). SNRs are measured by computing the mean source flux within a cylinder around the source position with a radius equal to the beam size and a height equal to the line width (source cylinder). The noise RMS is computed within an a hollow cylinder with the same center and volume of the source cylinder but with an inner radius starting at tree times the beam size.

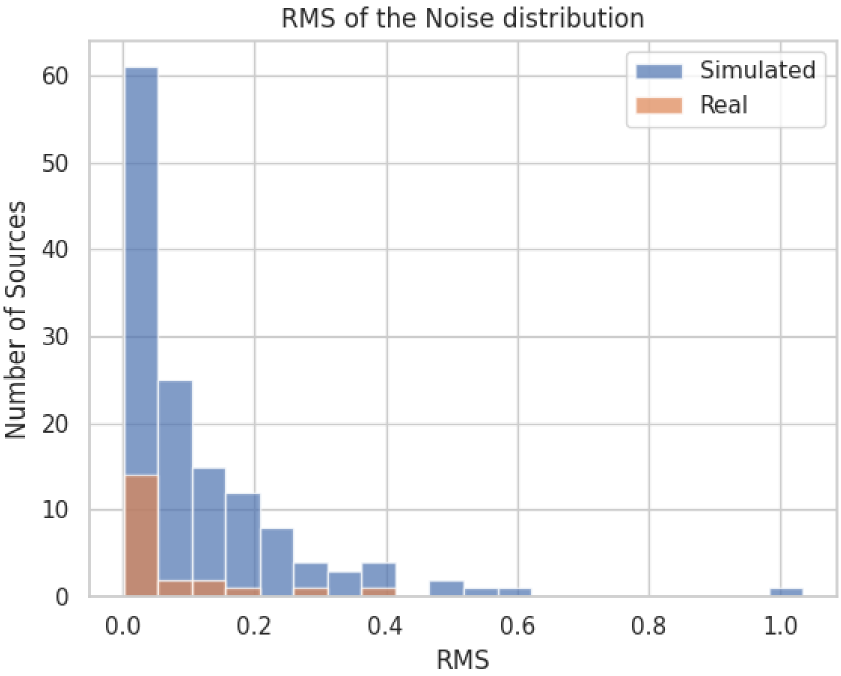

Figure 47: Comparison of the distributions of real (orange) and simulated (blue) noise RMSs. RMSs are measured by randomly sampling beam sizes chunks away from sources.

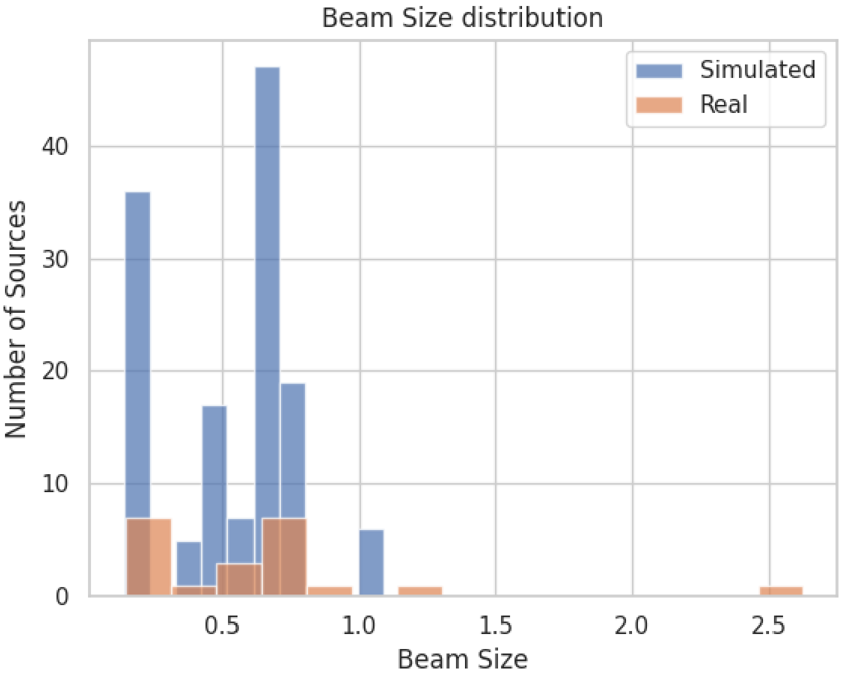

Figure 48: Comparison of the distributions of real (orange) and simulated (blue) Beam Sizes.

- **Isolating high spatial frequency patterns** (HFP) from the FFT of the image: Noise patterns at scales smaller than the PSF (commonly generated by the largest baselines) are reproduced in the simulated image.
- **Isolating elliptical (PSF-shaped) and radial patterns**: These patterns are a consequence of the sparse coverage of the u-v plane.
- **Correlating simulated noise as in the original reference image**: To this purpose, a local autocorrelation function (ACF) is computed from the original image and used to replicate a similar correlation in the simulation.
- **Adding a real sky image to the simulated noise**: A user-provided real sky image can be added at a pre–selected significance level.
- **Reproducing the original pixel flux distribution of each noise component**.

As shown in Figure 50, `NOISEMPIRE` can effectively identify and replicate noise patterns from real images in simulations. Being capable of detecting patterns at different scales, detecting issues on specific baselines during data acquisition will be enabled, allowing for timely interventions.

The current `ALMAsim` version only includes a beta version of the `NOISEMPIRE` module. The full integration of `NOISEMPIRE` into the `ALMAsim` package will enable to produce purely empirical, purely theoretical or mixed simulations. In addition, training AI algorithms on `NOISEMPIRE`-identified real patterns is expected to produce more realistic results. In the interim, `NOISEMPIRE` is being developed as a standalone tool. This two-step approach has been adopted for multiple reasons: i) it allows for a modular approach, making the code easier to understand, maintain, and test; ii) it minimizes the risk of delays or complete blocks during the development phase; iii) it allows for parallel development (by different people); iv) it maximizes the return of our development study,

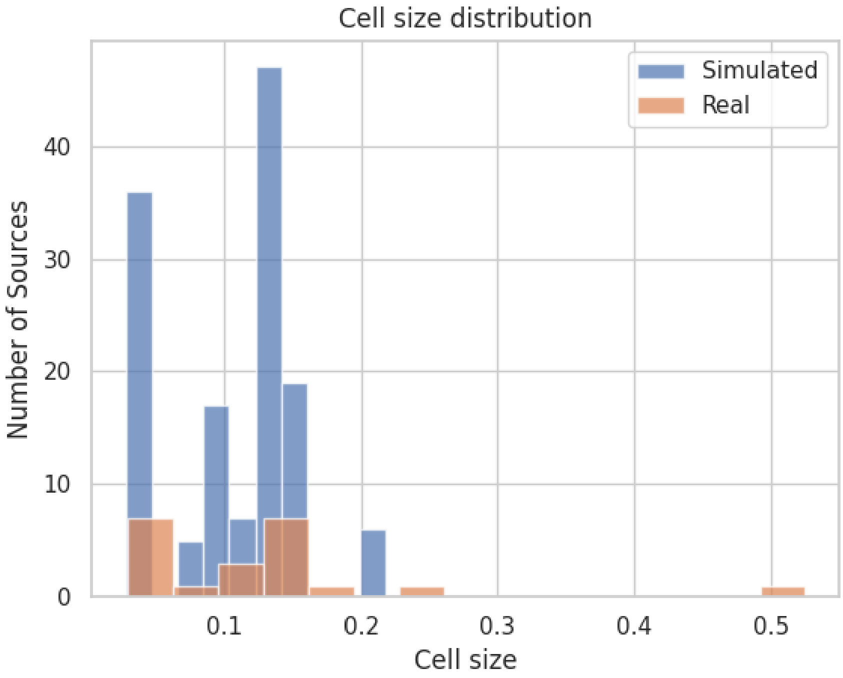

Figure 49: Comparison of the distributions of real (orange) and simulated (blue) Cell sizes or pixel dimensions in arcseconds.

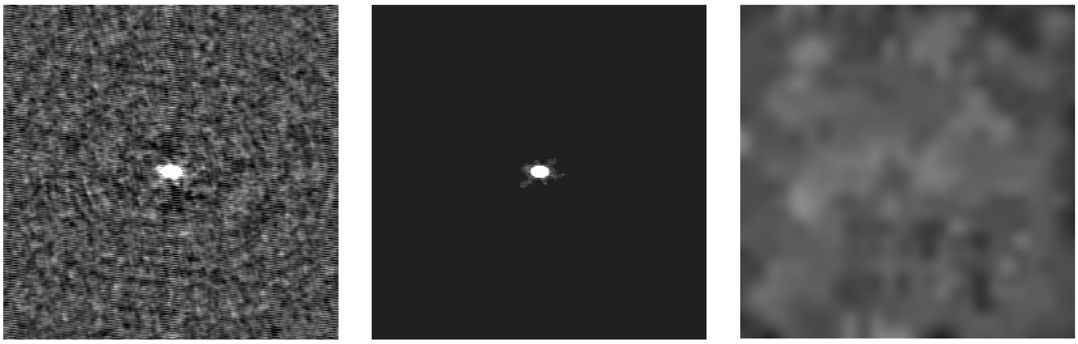

Original image (calibr. J1427-4206

Sources in the original image

Local underlying RMS

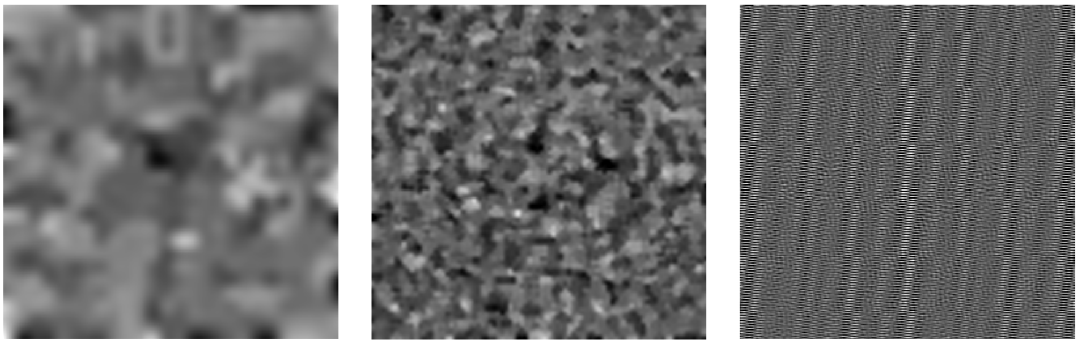

Large scale background

Small scale background

High spatial frequency patterns

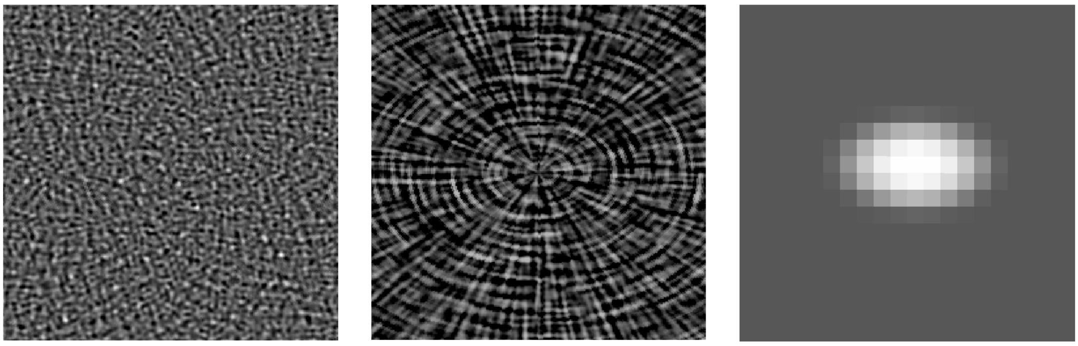

Noise correlation pattern

Radial and elliptical patterns

PSF in the original image

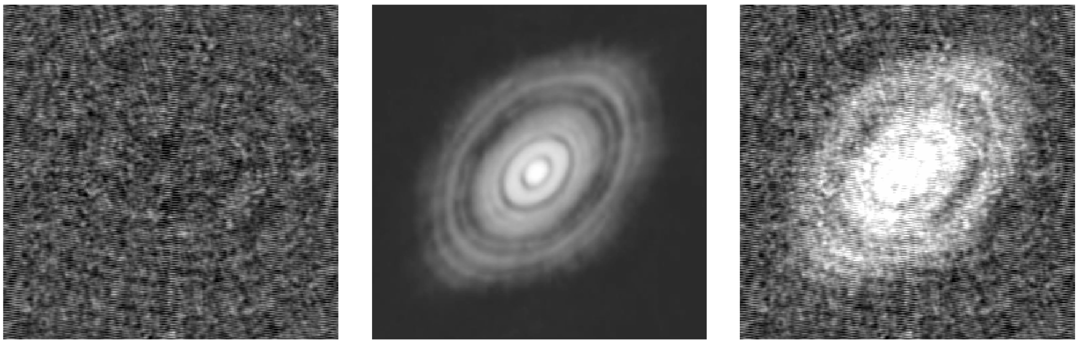

Simulated image (noise only)

Simulated real sky (HL-Tau protop. disk)

Simulated observation

Figure 50: Examples of noise components, features and patterns isolated by our prototype (`NOISEMPIRE` v1.0.2) from a real ALMA band 3 image of calibrator J1427-4206 (upper left panel). The same noise structure and patterns are used to simulate a pure noise image (bottom left panel) and an image of the HL-Tau protoplanetary disk (bottom central panel) as observed under similar conditions (bottom right panel).

as `NOISEMPIRE` can be readily downloaded as is, and used with images from various observatories, including different telescopes and wavelengths; v) Thanks to its ability to detect patterns, in the near future `NOISEMPIRE` can also be employed as an alarm system during ALMA operations, for example, to detect corrupted baselines in real time.

8 Final Remarks

We showcased the potential of two AI-based techniques to address challenges in ALMA synthesis imaging.

RESOLVE excels at detecting both diffuse emissions and point sources and will outperform other methods in combining data from multiple MSs for group-level imaging. Future steps include refining cube imaging, group imaging, self-calibration, full polarization imaging, and joint analyses with other ALMA projects and observatories, making it a leading tool within the scientific community.

DeepFocus, a meta-learner, significantly accelerates image analysis and is poised to become the solution of choice for real-time imaging, offering rapid support for operations like AoD tasks and enhancing archive completeness.

The companion tool, **ALMASim**, supports machine learning development by benchmarking performance on both real and simulated data. It improves ALMA dataset analysis and is refining ALMA simulations through **NOISEMPIRE**. These efforts include simulating diverse source types and integrating physical models like galaxy dynamics for advanced analysis.

Aligned with the **ALMA 2030 Vision**, our objectives include extending software functionalities, expanding applications to both real and simulated ALMA datasets, and publishing results in high-impact scientific journals to reach a broad international audience. Additionally, we aim to address key challenges in user interpretation, facilitating broader adoption of the tools.

Furthermore, we plan to enhance QA0-level operations through AQUA-WSU by implementing real-time image analysis for quick previews following calibration. Automated alerts will notify system astronomers via PRTSPR tickets to enable prompt responses. These advancements will integrate seamlessly with the Science Platform, providing real-time access to MOUS, improving workflows, and supporting joint analyses with observatories such as JWST.

8.1 Integration into CASA and RADPS

Considering feedback from the panel, we are exploring avenues for integrating these tools into Radio AStronomy Data Processing System (RADPS) or the legacy CASA framework, keeping in mind that the latter will be retired in the future. RADPS, being designed for compatibility with externally contributed software, offers potential for collaboration.

From a **user perspective**, `RESOLVE` and `DeepFocus` demonstrate substantial benefits when employed outside the constraints of major-minor `CASA` Cycles. Their independent advancements highlight the advantages of operating as standalone tools.

DeepFocus would not align well with CASA standalone due to mismatches in Python versions and technologies. However, officially entering the CASA/RADPS collaboration would be beneficial, particularly for furthering the development of `DeepFocus` and `ALMASim`. The CASA team's resources and expertise could accelerate their integration and functionality.

For **RESOLVE**, remaining outside `CASA` offers greater development flexibility. While independence allows rapid innovation, integrating specific capabilities into the next-generation `CASA` (ngCASA) team, i.e. RADPS, could foster shared advancements.

We remain open to discussions with NRAO's Research and Development group and the RADPS team. Collaborations with RADPS are possible if common ground is established; however, this depends on further defining and clarifying RADPS's scope, even though it is intended to align with ALMA's evolving scientific and operational objectives within the WSU framework.

Together, `RESOLVE`, `DeepFocus`, and `ALMASim` with `NOISEMPIRE` will drive advancements in data mining, analysis, and operational support, contributing significantly to the future of ALMA and its scientific community.

9 Appendix

- A BRAIN Study to Tackle Image Analysis with Artificial Intelligence in the ALMA 2030 Era (Guglielmetti, F. et al.)
- Bayesian and Machine Learning Methods in the Big Data era for astronomical imaging (Guglielmetti, F. et al.)
- Bayesian statistics approach to imaging of aperture synthesis data: RESOLVE meets ALMA (Tychoniec, L. et al.)
- 3D Detection and Characterisation of ALMA Sources through Deep Learning (Delli Veneri, M. et al.)
- Software:
    - ALMASim (Delli Veneri, M.)
    - DeepFocus (Delli Veneri, M.)
    - RESOLVE (IFT Group at MPA)
    - NIFTy (IFT Group at MPA)
    - NOISEmpire (Baronchelli, I.)

References

[1] H. Junklewitz, M. R. Bell, M. Selig, and T. A. Enßlin. RESOLVE: A new algorithm for aperture synthesis imaging of extended emission in radio astronomy. *A&A*, 586, Feb. 2016.

[2] M. Delli Veneri, Ł. Tychoniec, F. Guglielmetti, G. Longo, and E. Villard. 3D detection and characterization of ALMA sources through deep learning. *MNRAS*, 518(3):3407–3427, January 2023.

[3] J. Carpenter, D. Iono, F. Kemper, and A. Wootten. The alma development program: Roadmap to 2030. (2020).

[4] A. Gonzalez, P. Jewell, E. Humphreys, M. Zwaan, M. Fukagawa, N. Mizuno, J. Carpenter, T.-C. Shen, and S. Dougherty. The ALMA2030 wideband sensitivity upgrade. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, page 131022C. International Society for Optics and Photonics, SPIE, 2024.

[5] M. Delli Veneri. Deepfocus. `https://github.com/MicheleDelliVeneri/DeepFocus`, 2022/ 2024. Accessed: 09.09.2024.

[6] J. Roth. resolve. `https://gitlab.mpcdf.mpg.de/ift/resolve`, 2019/2024. Accessed: 09.09.2024.

[7] J. Huang, S. M. Andrews, C. P. Dullemond, A. Isella, L. M. Pérez, V. V. Guzmán, K. I. Öberg, Z. Zhu, S. Zhang, X.-N. Bai, M. Benisty, T. Birnstiel, J. M. Carpenter, A. M. Hughes, L. Ricci, E. Weaver, and D. J. Wilner. The disk substructures at high angular resolution project (dsharp). ii. characteristics of annular substructures. *ApJL*, 869(2):L42, December 2018.

[8] S. Martín, J. G. Mangum, N. Harada N., F. Costagliola, K. Sakamoto, S. Muller, R. Aladro, K. Tanaka, Y. Yoshimura, K. Nakanishi, R. Herrero-Illana, S. Mühle, S. Aalto, E. Behrens, L. Colzi, K. L. Emig, G. A. Fuller, S. García-Burillo, T. R. Greve, C. Henkel, J. Holdship, P. Humire, L. Hunt, T. Izumi, K. Kohno, S. König, D. S. Meier, T. Nakajima, Y. Nishimura, M. Padovani, V. M. Rivilla, S. Takano, P. P. van der Werf, S. Viti, and Y. T. Yan. ALCHEMI, an ALMA Comprehensive High-resolution Extragalactic Molecular Inventory. Survey presentation and first results from the ACA array. *A&A*, 656:A46, December 2021.

[9] M. Delli Veneri. Almasim. `https://github.com/MicheleDelliVeneri/ALMASim`, 2023/ 2024. Accessed: 09.09.2024.

[10] P. Frank. Nifty_7. `https://gitlab.mpcdf.mpg.de/ift/nifty/-/tree/NIFTy_7`, 2019/2024. Accessed: 09.09.2024.

[11] The CASA Team et al. CASA documentation. `https://casadocs.readthedocs.io/en/stable/`. Accessed on 05-09-2024.

[12] J. A. Hogbom. Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. *A&AS*, 15:417, June 1974.

[13] J. Roth, P. Frank, H. L. Bester, O. M. Smirnov, R. Westermann, and T. A. Enßlin. fast-resolve: Fast bayesian radio interferometric imaging. *A&A*, 690:A387, 2024.

[14] S. Sekhar, U. Rau, and R. Xue. Cube parallelization with casa. CASA Memo 13, National Radio Astronomy Observatory, August 2024.

[15] J. Roth, P. Arras, M. Reinecke, R. A. Perley, R. Westermann, and T. A. Enßlin. Bayesian radio interferometric imaging with direction-dependent calibration. *A&A*, 678:A177, 2023.

[16] J.-S. Kim, A. S. Nikonov, J. Roth, T. A. Enßlin, M. Janssen, P. Arras, H. Müller, and A. P. Lobanov. Bayesian self-calibration and imaging in very long baseline interferometry. *A&A*, 690:A129, 2024.

[17] J.-S. Kim, H. Müller, A. S. Nikonov, R.-S. Lu, J. Knollmüller, T. A. Enßlin, M. Wielgus, and A. P. Lobanov. Imaging the black hole shadow and extended jet of M87. *arXiv preprint arXiv:2409.00540*, 2024.

[18] P. A. Arras. Radio interferometry with information field theory, Mar. 2021.

[19] I. Baronchelli. NOISEmpire. `https://github.com/Ibaronch/NOISEMPIRE`, 2024. Accessed: 09.09.2024.

[20] The CASA Team, B. Bean, S. Bhatnagar, S. Castro, J. D. Meyer, B. Emonts, E. Garcia, R. Garwood, K. Golap, J. G. Villalba, P. Harris, Y. Hayashi, J. Hoskins, M. Hsieh, P. Jagannathan, W. Kawasaki, A. Keimpema, M. Kettenis, J. Lopez, J. Marvil, J. Masters, A. McNichols, D. Mehringer, R. Miel, G. Moellenbrock, F. Montesino, T. Nakazato, J. Ott, D. Petry, M. Pokorny, R. Raba, U. Rau, D. Schiebel, N. Schweighart, S. Sekhar, K. Shimada, D. Small, J.-W. Steeb, K. Sugimoto, V. Suoranta, T. Tsutsumi, I. M. van Bemmel, M. Verkouter, A. Wells, W. Xiong, A. Szomoru, M. Griffith, B. Glendenning, and J. Kern. Casa, the common astronomy software applications for radio astronomy. *PASP*, 134(1041):114501, November 2022.

[21] D. Nelson, A. Pillepich, V. Springel, R. Weinberger, L. Hernquist, R. Pakmor, S. Genel, P. Torrey, M. Vogelsberger, G. Kauffmann, et al. The IllustrisTNG simulations: Public data release. *Comput. Astrophys. Cosmol.*, 6(1):1–17, 2019.

[22] F. Guglielmetti, P. Arras, M. Delli Veneri, T. A. Enßlin, G. Longo, Ł. Tychoniec, and E. Villard. Bayesian and machine learning methods in the big data era for astronomical imaging. *Phys. Sci. Forum*, 5(1), 2022.

[23] Ł. Tychoniec, F. Guglielmetti, P. Arras, T. Enßlin, and E. Villard. Bayesian Statistics Approach to Imaging of Aperture Synthesis Data: RESOLVE Meets ALMA. *Phy. Sci. Forum*, 5(1), 2022.

[24] F. Guglielmetti, A. Mérand, M. Wittkowski, Ł. Tychoniec, G. González-Torà, M. A. Zwaan, P. Andreani, and C. De Breuck. Report on the ESO workshop: VLTI and ALMA Synthesis Imaging Workshop. *The Messenger*, 191:50–55, September 2023.

[25] F. Guglielmetti, M. Delli Veneri, I. Baronchelli, C. Blanco, A. Dosi, T. A. Enßlin, V. Johnson, G. Longo, J. Roth, F. Stoehr, Ł. Tychoniec, and E. Villard. A BRAIN Study to Tackle Image Analysis with Artificial Intelligence in the ALMA 2030 Era. *Phy. Sci. Forum*, 9(1), 2023.

[26] F. Guglielmetti. A Brain Study to Tackle Imaging with Artificial Intelligence in the ALMA 2030 ERA. *Video Memorie della Società Astronomica Italiana*, 4:29, December 2023.

[27] M. Delli Veneri. Deep Focus and ALMASim. *Video Memorie della Società Astronomica Italiana*, 4:4, December 2023.

[28] F. Guglielmetti, M. Delli Veneri, I. Baronchelli, V. Johnson, P. Arras, C. Blanco, M. Brescia, A. Dosi, T. A. Enßlin, G. Longo, J. Roth, F. Stoehr, Ł. Tychoniec, and E. Villard. A BRAIN study to tackle imaging with artificial intelligence in the ALMA2030 era. In *ALMA at 10 years: Past, Present, and Future*, page 26, December 2023.

[29] F. Guglielmetti. A BRAIN study to tackle imaging in the ALMA 2030 era. In *The Promises and Challenges of the ALMA Wideband Sensitivity Upgrade*, page 15, September 2024.

[30] E. D. Feigelson and G. J. Babu. *Modern Statistical Methods for Astronomy: With R Applications*. Cambridge University Press, 2012.

[31] A. Siemiginowska, G. Eadie, I. Czekala, E. Feigelson, E. B. Ford, V. Kashyap, M. Kuhn, T. Loredo, M. Ntampaka, A. Stevens, A. Avelino, K. Borne, T. Budavari, B. Burkhart, J. Cisewski-Kehe, F. Civano, I. Chilingarian, D. A. van Dyk, G. Fabbiano, D. P. Finkbeiner, D. Foreman-Mackey, P. Freeman, A. Fruscione, A. A. Goodman, M. Graham, H. M. Guenther, J. Hakkila, L. Hernquist, D. Huppenkothen, D. J. James, C. Law, J. Lazio, T. Lee, M. López-Morales, A. A. Mahabal, K. Mandel, X.-L. Meng, J. Moustakas, D. Muna, J. E. G. Peek, G. Richards, S. K. N. Portillo, J. Scargle, R. S. de Souza, J. S. Speagle, K. G. Stassun, D. C. Stenning, S. R. Taylor, G. R. Tremblay, V. Trimble, P. A. Yanamandra-Fisher, and C. A. Young. The Next Decade of Astroinformatics and Astrostatistics. *BAAS*, 51(3):355, May 2019.

[32] G. Longo, M. Brescia, S. G. Djorgovski, S. Cavuoti, and C. Donalek. Data Driven Discovery in Astrophysics. *arXiv preprint arXiv:1410.5631*, 2014.

[33] K. D. Borne, R. Ghani, R. Murphy, and W. S. Jacoby. The Emerging Field of Astroinformatics. *arXiv preprint arXiv:0911.0505*, 2009.

[34] J. Rüstig, M. Guardiani, J. Roth, P. Frank, and T. Enßlin. Introducing lenscharm - a charming bayesian strong lensing reconstruction framework. *A&A*, 682:A146, 2024.

[35] L. I. Scheel-Platz, J. Knollmüller, P. Arras, P. Frank, M. Reinecke, D. Jüstel, and T. A. Enßlin. Multicomponent imaging of the Fermi gamma-ray sky in the spatio-spectral domain. *A&A*, 680:A2, 2023.

[36] N. Reeb, S. Hutschenreuter, P. Zehetner, T. Ensslin, S. Alves, M. André, M. Anghinolfi, G. Anton, M. Ardid, J. J. Aubert, J. Aublin, B. Baret, S. Basa, B. Belhorma, M. Bendahman, V. Bertin, S. Biagi, M. Bissinger, J. Boumaaza, M. Bouta, M. C. Bouwhuis, H. Brânzaş, R. Bruijn, J. Brunner, J. Busto, B. Caiffi, A. Capone, L. Caramete, J. Carr, V. Carretero, S. Celli, M. Chabab, T. N. Chau, R. Cherkaoui El Moursli, T. Chiarusi, M. Circella, A. Coleiro, M. Colomer-Molla, R. Coniglione, P. Coyle, A. Creusot,

A. F. Díaz, G. de Wasseige, A. Deschamps, C. Distefano, I. Di Palma, A. Domi, C. Donzaud, D. Dornic, D. Drouhin, T. Eberl, T. van Eeden, N. El Khayati, A. Enzenhöfer, P. Fermani, G. Ferrara, F. Filippini, L. Fusco, Y. Gatelet, P. Gay, H. Glotin, R. Gozzini, R. Gracia Ruiz, K. Graf, C. Guidi, S. Hallmann, H. van Haren, A. J. Heijboer, Y. Hello, J. J. Hernández-Rey, J. Hößl, J. Hofestädt, F. Huang, G. Illuminati, C. W. James, B. Jisse-Jung, M. d. Jong, P. d. Jong, M. Jongen, M. Kadler, O. Kalekin, U. Katz, N. R. Khan-Chowdhury, A. Kouchner, I. Kreykenbohm, V. Kulikovskiy, R. Lahmann, R. Le Breton, D. Lefèvre, E. Leonora, G. Levi, M. Lincetto, D. Lopez-Coto, S. Loucatos, L. Maderer, J. Manczak, M. Marcelin, A. Margiotta, A. Marinelli, J. A. Martínez-Mora, K. Melis, P. Migliozzi, A. Moussa, R. Muller, L. Nauta, S. Navas, E. Nezri, B. Ó Fearraigh, M. Organokov, G. E. Păvălaş, C. Pellegrino, M. Perrin-Terrin, P. Piattelli, C. Pieterse, C. Poirè, V. Popa, T. Pradier, N. Randazzo, S. Reck, G. Riccobene, A. Romanov, A. Sánchez-Losa, F. Salesa Greus, D. F. E. Samtleben, M. Sanguineti, P. Sapienza, J. Schnabel, J. Schumann, F. Schüssler, M. Spurio, Th. Stolarczyk, M. Taiuti, Y. Tayalati, S. J. Tingay, B. Vallage, V. Van Elewyck, F. Versari, S. Viola, D. Vivolo, J. Wilms, S. Zavatarelli, A. Zegarelli, J. D. Zornoza, and J. Zúñiga. Studying Bioluminescence Flashes with the ANTARES Deep Sea Neutrino Telescope. *arXiv e-prints arXiv:2107.08063*, July 2021.

[37] M. Guardiani, P. Frank, A. Kostić, G. Edenhofer, J. Roth, B. Uhlmann, and T. Enßlin. Causal, Bayesian, & non-parametric modeling of the SARS-CoV-2 viral load distribution vs. patient's age. *PLoS ONE*, 17(10):e0275011, October 2022.

[38] M. Brescia, S. Cavuoti, M. Paolillo, G. Longo, and T. Puzia. The detection of globular clusters in galaxies as a data mining problem. *MNRAS*, 421(2):1155–1165, 03 2012.

[39] D. Huppenkothen, M. Ntampaka, M. Ho, M. Fouesneau, B. Nord, J. E. G. Peek, M. Walmsley, J. F. Wu, C. Avestruz, T. Buck, M. Brescia, D. P. Finkbeiner, A. D. Goulding, T. Kacprzak, P. Melchior, M. Pasquato, N. Ramachandra, Y. S. Ting, G. van de Ven, S. Villar, V. A. Villar, and E. Zinger. Constructing Impactful Machine Learning Research for Astronomy: Best Practices for Researchers and Reviewers. *arXiv e-prints arXiv:2310.12528*, October 2023.

[40] M. Delli Veneri, S. Cavuoti, M. Brescia, G. Longo, and G. Riccio. Star formation rates for photometric samples of galaxies using machine learning methods. *MNRAS*, 486(1):1377–1391, 03 2019.

[41] R. D'Abrusco, G. Longo, and N. A. Walton. Quasar candidates selection in the Virtual Observatory era. *MNRAS*, 396(1):223–262, 06 2009.

[42] F. Guglielmetti, R. Fischer, and V. Dose. Background-source separation in astronomical images with bayesian probability theory - i. the method. *MNRAS*, 396:165–190, 2009.

[43] K-E Harabi, T. Hirtzlin, T. Clément, E. Vianello, R. Laurent, J. Droulez, P. Bessière, J-M Portal, M. Bocquet, and D. Querlioz. A memristor-based Bayesian machine. *Nature Electronics*, 6:52–63, December 2022.

[44] F. Guglielmetti, E. Villard, and E. Fomalont. Bayesian reconstruction through adaptive image notion. In *Bayesian Inference and Maximum Entropy Methods in Science and Engineering—MaxEnt 2019*, volume 33, 8, pages 95–101, 30 June - 5 July 2019.

[45] T. J. Cornwell and K. F. Evans. A simple maximum entropy deconvolution algorithm. *A&A*, 143:77–83, February 1985.

[46] A. W. Strong. Maximum Entropy Imaging of COMPTEL Data. *Exp. Astron.*, 6(4):97–102, Dec. 1995.

[47] M. Cárcamo, P.E. Román, S. Casassus, V. Moral, and F.R. Rannou. Multi-gpu maximum entropy image synthesis for radio astronomy. *Astron. Comput.*, 22:16–27, 2018.

[48] A. Levis, A. A. Chael, K. L. Bouman, M. Wielgus, and P. P. Srinivasan. Orbital polarimetric tomography of a flare near the sagittarius a* supermassive black hole. *Nature Astronomy*, 8:765–773, apr 2024. Open access, published 22 April 2024.

[49] E. Tolley, S. Frasch, E. Orliac, S. Krishna, M. Bianco, S. Kashani, P. Hurley, M. Simeoni, and J.-P. Kneib. BIPP: An efficient HPC implementation of the Bluebild algorithm for radio astronomy. *arXiv preprint arXiv:2310.09200*, 2024.

[50] B. N. Chiche, J. N. Girard, J. Frontera-Pons, A. Woiselle, and J.-L. Starck. Deep learning-based deconvolution for interferometric radio transient reconstruction. *A&A*, 675:A116, 2023.

[51] H. Eklund. Deep solar ALMA neural network estimator for image refinement and estimates of small-scale dynamics. *A&A*, 669:A106, 2023.

[52] B. Zawadzki, I. Czekala, R. A. Loomis, T. Quinn, H. Grzybowski, R. C. Frazier, J. Jennings, K. M. Nizam, and Y. Jian. Regularized Maximum Likelihood Image Synthesis and Validation for ALMA Continuum Observations of Protoplanetary Disks. *PASP*, 135(1048):064503, June 2023.

[53] C. Tasse, B. Hugo, M. Mirmont, O. Smirnov, M. Atemkeng, L. Bester, M. J. Hardcastle, R. Lakhoo, S. Perkins, and T. Shimwell. DDFacet: Facet-based radio imaging package. Astrophysics Source Code Library, record ascl:2305.008, May 2023.

[54] P. Tiede. Comrade: Composable Modeling of Radio Emission. *JOSS*, 7(76):4457, 2022.

[55] M. Yamaguchi, K. Akiyama, T. Tsukagoshi, T. Muto, A. Kataoka, F. Tazaki, S. Ikeda, M. Fukagawa, M. Honma, and R. Kawabe. Super-resolution Imaging of the Protoplanetary Disk HD 142527 Using Sparse Modeling. *ApJ*, 895(2):84, May 2020.

[56] C Gheller, F Vazza, and A Bonafede. Deep learning based detection of cosmological diffuse radio sources. *MNRAS*, 480(3):3749–3761, 08 2018.

[57] A. R. Offringa and O. Smirnov. An optimized algorithm for multiscale wideband deconvolution of radio astronomical images. *MNRAS*, 471(1):301–316, 06 2017.

[58] R. Vio, C. Vergès, and P. Andreani. The correct estimate of the probability of false detection of the matched filter in weak-signal detection problems - II. Further results with application to a set of ALMA and ATCA data. *A&A*, 604:A115, 2017.

[59] N. H. Hayatsu, Y. Matsuda, H. Umehata, N. Yoshida, I. Smail, A. M. Swinbank, R. Ivison, K. Kohno, Y. Tamura, M. Kubo, D. Iono, B. Hatsukade, K. Nakanishi, R. Kawabe, T. Nagao, A. K. Inoue, T. T. Takeuchi, M. Lee, Y. Ao, S. Fujimoto, T. Izumi, Y. Yamaguchi, S. Ikarashi, and T. Yamada. ALMA deep field in SSA22: Blindly detected CO emitters and [CII] emitter candidates. *PASJ*, 69(3):45, 05 2017.

[60] E. Merényi, J. Taylor, and A. Isella. Deep data: discovery and visualization application to hyperspectral alma imagery. *Proc. of the IAU*, 12(S325):281–290, 2016.

[61] P. Fagrell, A. Ybring, E. Östling, M. C. Toribio, J. Kainulainen, A. Plunkett, and P. Bjerkeli. Classifying ALMA continuum data using machine learning, November 2023.

[62] A. Emsenhuber, R. Burn, J. Weder, K. Monsch, G. Picogna, B. Ercolano, and T. Preibisch. Toward a population synthesis of disks and planets - II. Confronting disk models and observations at the population level. *A&A*, 673:A78, 2023.

[63] J. Jennings, M. Tazzari, C.J. Clarke, R. A. Booth, and G. P. Rosotti. Superresolution trends in the ALMA Taurus survey: structured inner discs and compact discs. *MNRAS*, 514(4):6053–6073, 06 2022.

[64] J. P. Terry, C. Hall, S. Abreau, and S. Gleyzer. Locating Hidden Exoplanets in ALMA Data Using Machine Learning. *ApJ*, 941(2):192, December 2022.

[65] I. Villicana-Pedraza, R. Walterbos, F. Carreto-Parra, J. Ott, E. Momjian, A. Thelen, A. Ginsburg, L. Zapata, M. Gonzalez, J. Floyd, J. Saucedo, L. Binette, and S. Prugh. Preliminary results from prebiotic molecules with ALMA in the era of artificial intelligence. In E. da Cunha, J. Hodge, J. Afonso, L. Pentericci, and D. Sobral, editors, *Uncovering Early Galaxy Evolution in the ALMA and JWST Era*, volume 352 of *IAU Symp.*, pages 248–250, January 2020.

[66] W. Von Der Linden, V. Dose, and U. Von Toussaint. *Bayesian Probability Theory: Applications in the Physical Sciences*. Wiley-VCH, Weinheim, 2009.

[67] T. A. Enßlin, M. Frommert, and F. S. Kitaura. Information field theory for cosmological perturbation reconstruction and nonlinear signal analysis. *Phys. Rev. D*, 80(10):105005, Nov. 2009.

[68] T. Enßlin. Information field theory. In U. von Toussaint, editor, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 32nd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 1553 of *American Institute of Physics Conference Series*, pages 184–191, Aug. 2013.

[69] M. Greiner, V. Vacca, H. Junklewitz, and T. A. Enßlin. fastRESOLVE: fast Bayesian imaging for aperture synthesis in radio astronomy. *arXiv e-prints*, page arXiv:1605.04317, May 2016.

[70] P. Arras, J. Knollmüller, H. Junklewitz, and T. A. Enßlin. Radio Imaging With Information Field Theory. *arXiv e-prints arXiv:1803.02174*, Mar. 2018.

[71] J. Knollmüller and T. A. Enßlin. Metric Gaussian Variational Inference. *arXiv preprint arXiv:1901.11033*, abs/1901.11033, 2019.

[72] P. Frank, R. Leike, and T. A. Enßlin. Geometric Variational Inference. *Entropy*, 23(7), 2021.

[73] J. Roth, M. Reinecke, and G. Edenhofer. Jaxbind: Bind any function to jax. *JOSS*, 9(98), 6532, 2024.

[74] C. L. Brogan and et al. THE 2014 ALMA LONG BASELINE CAMPAIGN: FIRST RESULTS FROM HIGH ANGULAR RESOLUTION OBSERVATIONS TOWARD THE HL TAU REGION. *ApJ*, 808(1):L3, July 2015.

[75] D. S. Briggs. High Fidelity Interferometric Imaging: Robust Weighting and NNLS Deconvolution. In *AAS Meeting Abstracts*, volume 187, page 112.02, Dec. 1995.

[76] P. Arras, H. L. Bester, R. A. Perley, R. Leike, O. Smirnov, R. Westermann, and T. A. Enßlin. Comparison of classical and Bayesian imaging in radio interferometry. *A&A*, 646:A84, Feb. 2021.

[77] J. Huang, S. M. Andrews, L. M. Perez, Z. Zhu, C. P. Dullemond, A. Isella, M. Benisty, X.-N. Bai, T. Birnstiel, J. M. Carpenter, V. V. Guzman, A. M. Hughes, K. I. Oberg, L. Ricci, D. J. Wilner, and S. Zhang. The Disk Substructures at High Angular Resolution Project (DSHARP). III. Spiral Structures in the Millimeter Continuum of the Elias 27, IM Lup, and WaOph 6 Disks. *ApJL*, (2), 2018.

[78] T. Mroczkowski, D. Nagai, P. Andreani, M. Arnaud, J. Bartlett, N. Battaglia, K. Basu, E. Bulbul, J. Chluba, E. Churazov, C. Cicone, A. Crites, M. Devlin, N. DeNigris, S. Dicker, L. Di Mascolo, M. Gaspari, S. Golwala, F. Guglielmetti, J. C. Hill, T. Kitayama, P. Klaassen, K. Kohno, E. Komatsu, R. Kneissl, M. Lacy, B. Mason, K. Nyland, C. Romero, J. Sayers, N. Sehgal, S. Simon, R. Sunyaev, G. Wilson, M. Zemcov, and J. ZuHone. A High-resolution SZ View of the Warm-Hot Universe. *Bulletin of the AAS*, 51(3), may 31 2019. https://baas.aas.org/pub/2020n3i124.

[79] L. Di Mascolo, Y. Perrott, T. Mroczkowski, S. Andreon, S. Ettori, A. Simionescu, S. Raghunathan, J. van Marrewijk, C. Cicone, M. Lee, D. Nelson, L. Sommovigo, M. Booth, P. Klaassen, P. Andreani, M. A. Cordiner, D. Johnstone, E. van Kampen, D. Liu, T. J. Maccarone, T. W. Morris, A. Saintonge, M. Smith, A. E. Thelen, and S. Wedemeyer. Atacama Large Aperture Submillimeter Telescope (AtLAST) Science: Resolving the Hot and Ionized Universe through the Sunyaev-Zeldovich effect. *arXiv preprint arXiv:2403.00909*, 2024.

[80] L. Di Mascolo, T. Mroczkowski, Y. Perrott, L. Rudnick, M. James Jee, K. HyeongHan, E. Churazov, J. D. Collier, J. M. Diego, A. M. Hopkins, J. Kim, B. S. Koribalski, J. D. Marvil, R. van der Burg, and J.-L. West. Multiwavelength view of SPT-CL J2106-5844 - The radio galaxies and the thermal and relativistic plasmas in a massive galaxy cluster merger at z 1.13. *A&A*, 650:A153, 2021.

[81] W. D. Cotton. Fourier Plane Image Combination by Feathering. *PASP*, 129(979):094501, July 2017.

[82] U. Rau, N. Naik, and T. Braun. A Joint Deconvolution Algorithm to Combine Single-dish and Interferometer Data for Wideband Multiterm and Mosaic Imaging. *AJ*, 158(1):3, June 2019.

[83] G. Edenhofer, C. Zucker, P. Frank, A. K. Saydjari, J. S. Speagle, D. Finkbeiner, and T. A. Enßlin. A parsec-scale Galactic 3D dust map out to 1.25 kpc from the Sun. *A&A*, 685:A82, May 2024.

[84] S. Carniani, A. Marconi, A. Biggs, G. Cresci, G. Cupani, V. D'Odorico, E. Humphreys, R. Maiolino, F. Mannucci, P. Molaro, T. Nagao, L. Testi, and M. A. Zwaan. Strongly star-forming rotating disks in a complex merging system at $z = 4.7$ as revealed by ALMA. *A&A*, 559:A29, Nov. 2013.

[85] J. P. McMullin, B. Waters, D. Schiebel, W. Young, and K. Golap. CASA Architecture and Applications. In R. A. Shaw, F. Hill, and D. J. Bell, editors, *ADASS XVI*, volume 376 of *Astronomical Society of the Pacific Conference Series*, page 127, Oct. 2007.

[86] R. Preuss and U. Von Toussaint. Global Optimization Employing Gaussian Process-Based Bayesian Surrogates. *Entropy*, 20(3), 2018.

[87] M. Delli Veneri, L. Sannino, A. Rownok, J. Roth, F. Guglielmetti, E. Villard, M. Brescia, and G. Longo. MicheleDelliVeneri/ALMASim: v2.1.9, July 2024.

[88] F. Marinacci, M. Vogelsberger, R. Pakmor, P. Torrey, V. Springel, L. Hernquist, D. Nelson, R. Weinberger, A. Pillepich, J. Naiman, and S. Genel. First results from the IllustrisTNG simulations: radio haloes and magnetic fields. *MNRAS*, 480(4):5113–5139, 08 2018.

[89] C. J. Lintott, K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu, et al. Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *MNRAS*, 389(3):1179–1189, 2008.

[90] Hubble Space Telescope. Hubble top 100 images. `https://hubblesite.org/collection/top-100`, 2020. Accessed: 2024-09-11.

[91] M. Rocklin. Dask: Parallel Computation with Blocked algorithms and Task Scheduling. `https://dask.org`, 2015. Accessed: 2024-09-11.

[92] A. B. Yoo, M. A. Jette, and M. Grondona. Slurm: Simple linux utility for resource management. In *Job Scheduling Strategies for Parallel Processing*, pages 44–60. Springer, 2003.

[93] M. Reinecke, T. Steininger, and M. Selig. NIFTy – Numerical Information Field TheorY.

[94] S. Andrianomena and H. Tang. Radio galaxy zoo: Leveraging latent space representations from variational autoencoder. *Journal of Cosmology and Astroparticle Physics*, 2024(06):034, jun 2024.

[95] S. T. Myers. The VLA Sky Survey (VLASS): First Half-Epoch (1.1) Results and Future Prospects. In *AAS Meet. Abs. #233*, volume 233 of *AAS Meet. Abs.*, page 214.03, January 2019.

[96] Robert H. Becker, Richard L. White, and David J. Helfand. The VLA's FIRST Survey. In D. R. Crabtree, R. J. Hanisch, and J. Barnes, editors, *Astronomical Data Analysis Software and Systems III*, volume 61 of *ASP Conf. Ser.*, page 165, January 1994.

[97] P. N. Best, R. Kondapally, W. L. Williams, R. K. Cochrane, K. J. Duncan, C. L. Hale, P. Haskell, K. Małek, I. McCheyne, D. J. B. Smith, L. Wang, A. Botteon, M. Bonato, M. Bondi, G. Calistro Rivera, F. Gao, G. Gürkan, M. J. Hardcastle, M. J. Jarvis, B. Mingo, H. Miraghaei, L. K. Morabito, D. Nisbet, I. Prandoni, H. J. A. Röttgering, J. Sabater, T. Shimwell, C. Tasse, and R. van Weeren. The LOFAR Two-metre Sky Survey: Deep Fields data release 1. V. Survey description, source classifications, and host galaxy properties. *MNRAS*, 523(2):1729–1755, August 2023.

[98] Heald, G. H., Pizzo, R. F., Orrú, E., Breton, R. P., Carbone, D., Ferrari, C., Hardcastle, M. J., Jurusik, W., Macario, G., Mulcahy, D., Rafferty, D., Asgekar, A., Brentjens, M., Fallows, R. A., Frieswijk, W., Toribio, M. C., Adebahr, B., Arts, M., Bell, M. R., Bonafede, A., Bray, J., Broderick, J., Cantwell, T., Carroll, P., Cendes, Y., Clarke, A. O., Croston, J., Daiboo, S., de Gasperin, F., Gregson, J., Harwood, J., Hassall, T., Heesen, V., Horneffer, A., van der Horst, A. J., Iacobelli, M., Jelić, V., Jones, D., Kant, D., Kokotanekov, G., Martin, P., McKean, J. P., Morabito, L. K., Nikiel-Wroczyński, B., Offringa, A., Pandey, V. N., Pandey-Pommier, M., Pietka, M., Pratley, L., Riseley, C., Rowlinson, A., Sabater, J., Scaife, A. M. M., Scheers, L. H. A., Sendlinger, K., Shulevski, A., Sipior, M., Sobey, C., Stewart, A. J., Stroe, A., Swinbank, J., Tasse, C., Trüstedt, J., Varenius, E., van Velzen, S., Vilchez, N., van Weeren, R. J., Wijnholds, S., Williams, W. L., de Bruyn, A. G., Nijboer, R., Wise, M., Alexov, A., Anderson, J., Avruch, I. M., Beck, R., Bell, M. E., van Bemmel, I., Bentum, M. J., Bernardi, G., Best, P., Breitling, F., Brouw, W. N., Brüggen, M., Butcher, H. R., Ciardi, B., Conway, J. E., de Geus, E., de Jong, A., de Vos, M., Deller, A., Dettmar, R.-J., Duscha, S., Eislöffel, J., Engels, D., Falcke, H., Fender, R., Garrett, M. A., Grießmeier, J., Gunst, A. W., Hamaker, J. P., Hessels, J. W. T., Hoeft, M., Hörandel, J., Holties, H. A., Intema, H., Jackson, N. J., Jütte, E., Karastergiou, A., Klijn, W. F. A., Kondratiev, V. I., Koopmans, L. V. E., Kuniyoshi, M., Kuper, G., Law, C., van Leeuwen, J., Loose, M., Maat, P., Markoff, S., McFadden, R., McKay-Bukowski, D., Mevius, M., Miller-Jones, J. C. A., Morganti, R., Munk, H., Nelles, A., Noordam, J. E., Norden, M. J., Paas, H., Polatidis, A. G., Reich, W., Renting, A., Röttgering, H., Schoenmakers, A., Schwarz, D., Sluman, J., Smirnov, O., Stappers, B. W., Steinmetz, M., Tagger, M., Tang, Y., ter Veen, S., Thoudam, S., Vermeulen, R., Vocks, C., Vogt, C., Wijers, R. A. M. J., Wucknitz, O., Yatawatta, S., and Zarka, P. The lofar multifrequency snapshot sky survey (msss) - i. survey description and first results. *A&A*, 582:A123, 2015.

[99] Z.-Y. Cai, A. Lapi, J.-Q. Xia, G. De Zotti, M. Negrello, C. Gruppioni, E. Rigby, and et al. The Evolution of Far-infrared and Radio Luminosity Correlations in Star-forming Galaxies. *ApJ*, 768:21, 2013.

[100] J. S. Spilker, D. P. Marrone, J. E. Aguirre, M. Aravena, M. L. N. Ashby, M. Béthermin, C. M. Bradford, M. S. Bothwell, M. Brodwin, J. E. Carlstrom, et al. The rest-frame submillimeter spectrum of high-redshift, dusty, star-forming galaxies. *ApJ*, 785(2):149, 2014.

[101] A. Kirkpatrick, C. Sharon, E. Keller, and A. Pope. CO Emission in Infrared-selected Active Galactic Nuclei. *ApJ*, 879(1):41, 2019.

[102] B. P. Venemans, F. Walter, M. Neeleman, M. Novak, J. Otter, R. Decarli, E. Bañados, A. Drake, E. P. Farina, M. Kaasinen, C. Mazzucchelli, C. Carilli, X. Fan, H.-W. Rix, and R. Wang. Kiloparsec-scale ALMA Imaging of [C II] and Dust Continuum Emission of 27 Quasar Host Galaxies at $z \sim 6$. *ApJ*, 904(2):130, 2020.

[103] T. Mroczkowski, P. A. Gallardo, M. Timpe, A. Kiselev, M. Groh, H. Kaercher, M. Reichert, C. Cicone, R. Puddu, P. Dubois dit Bonclaude, D. Bok, E. Dahl, M. Macintosh, S. Dicker, I. Viole, S. Sartori, G. A. Valenzuela Venegas, M. Zeyringer, M. Niemack, S. Poppi, R. Olguin, E. Hatziminaoglou, C. De Breuck, P. Klaassen, F. M. Montenegro-Montes, and T. Zimmerer. The conceptual design of the 50-meter atacama large aperture submillimeter telescope (atlast). *arXiv preprint arXiv:2402.18645*, 2024.

[104] J. van Marrewijk, T. W. Morris, T. Mroczkowski, C. Cicone, S. Dicker, L. Di Mascolo, S. K. Haridas, J. Orlowski-Scherer, E. Rasia, and C. Romero. maria: A novel simulator for forecasting (sub-)mm observations. *arXiv preprint arXiv:2402.10731*, 2024.

[105] I. Baronchelli, M. Bonato, G. De Zotti, V. Casasola, M. Delli Veneri, F. Guglielmetti, E. Liuzzo, R. Paladino, L. Trobbiani, and M. Zwaan. ALMA Band 3 Source Counts: A Machine Learning Approach to Contamination Mitigation below 5 Sigma. *Galaxies*, 12(3):26, May 2024.

[106] E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *A&AS*, 117:393–404, June 1996.