

R2D2 image reconstruction with model uncertainty quantification in radio astronomy

Amir Aghabiglou¹, Chung San Chu¹, Arwa Dabbech¹, Yves Wiaux^{1†}

¹Institute of Sensors, Signals and Systems, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom

Email: [†]y.wiaux@hw.ac.uk

Abstract—The “Residual-to-Residual DNN series for high-Dynamic range imaging” (R2D2) approach was recently introduced for Radio-Interferometric (RI) imaging in astronomy. R2D2’s reconstruction is formed as a series of residual images, iteratively estimated as outputs of Deep Neural Networks (DNNs) taking the previous iteration’s image estimate and associated data residual as inputs. In this work, we investigate the robustness of the R2D2 image estimation process, by studying the uncertainty associated with its series of learned models. Adopting an ensemble averaging approach, multiple series can be trained, arising from different random DNN initializations of the training process at each iteration. The resulting multiple R2D2 instances can also be leveraged to generate “R2D2 samples”, from which empirical mean and standard deviation endow the algorithm with a joint estimation and uncertainty quantification functionality. Focusing on RI imaging, and adopting a telescope-specific approach, multiple R2D2 instances were trained to encompass the most general observation setting of the Very Large Array (VLA). Simulations and real-data experiments confirm that: (i) R2D2’s image estimation capability is superior to that of the state-of-the-art algorithms; (ii) its ultra-fast reconstruction capability (arising from series with only few DNNs) makes the computation of multiple reconstruction samples and of uncertainty maps practical even at large image dimension; (iii) it is characterized by a very low model uncertainty.

Index Terms—Computational methods, Deep learning, Astronomy image processing, Aperture synthesis

I. INTRODUCTION

In the field of imaging sciences, one of the significant challenges is to address ill-posed imaging inverse problems while ensuring accurate reconstructions and scalability to large datasets. Aperture synthesis by RI in astronomy is an iconic example of this challenge towards joint precision and scalability. In RI, each pair of antennas captures noisy complex measurements, termed visibilities, corresponding to Fourier components of the sought radio emission [1]. These measurements, gathered over the observation period, form an incomplete coverage of the 2D Fourier plane, constituting the Fourier sampling pattern.

Consider $x^* \in \mathbb{R}_+^N$ as the unknown radio image of interest and focusing on monochromatic intensity imaging, the observed data $y \in \mathbb{C}^M$ can be expressed as:

$$y = \Phi x^* + n, \quad (1)$$

where $n \in \mathbb{C}^M$ represents complex random noise following a Gaussian distribution $\mathcal{N}(0, \tau^2)$, with $\tau > 0$. Here, $\Phi: \mathbb{R}^N \rightarrow \mathbb{C}^M$ is the measurement operator, describing the non-uniform Fourier sampling [1] and often incorporating a data-weighting scheme to improve the observation’s effective resolution [2].

RI data can be formulated in the image domain via a back-projection, resulting in the *dirty* image $x_d \in \mathbb{R}^N$ given by:

$$x_d = \kappa \text{Re}\{\Phi^\dagger y\} = \kappa \text{Re}\{\Phi^\dagger \Phi x^* + \Phi^\dagger n\}, \quad (2)$$

where $\kappa > 0$ serves as a normalization factor, ensuring that the point spread function (PSF) has a peak value of one. The linear operator $\kappa \text{Re}\{\Phi^\dagger \Phi\}: \mathbb{R}^N \rightarrow \mathbb{R}^N$ maps the image of interest to the dirty image space, with $\text{Re}\{\cdot\}$ indicating the real part of this dirty image. The term $\kappa \text{Re}\{\Phi^\dagger n\}$ represents the image-domain noise vector.

The CLEAN algorithm [3], a bespoke version of the Matching Pursuit algorithm [4] that iteratively identifies model components from residual dirty images, holds a near-total monopoly in the field. While its simplicity and computational efficiency endows the algorithm with a high degree of scalability, it also limits its capability in terms of reconstruction quality. Over the past decade, new RI imaging algorithms have emerged, aiming to supersede CLEAN. Optimization algorithms propelled by advanced sparsity-based regularization models have been proposed [5]–[9], with uSARA the latest evolution [10]. Recent advances in deep learning have opened new avenues in RI imaging, offering both modeling power and acceleration. These range from end-to-end approaches such as POLISH [10]–[12], to Plug-and-Play (PnP) algorithms such as AIRI [10], whose regularization operators are encapsulated in DNN denoisers. While such optimization and PnP algorithms have demonstrated remarkable precision indeed, their highly iterative nature poses challenges towards scaling to the extreme data volumes expected from modern telescopes such as the VLA¹ and its next-generation upgrade ngVLA.

In recent work, a novel deep learning-based algorithm was proposed for RI imaging, termed “Residual-to-Residual DNN series for high-Dynamic range imaging” (R2D2) [13], [14]. R2D2’s reconstruction is formed as a series of residual images, iteratively estimated as outputs of regularization DNNs taking the previous iteration’s image estimate and associated data residual as inputs. The method can be interpreted as a learned version of Matching Pursuit. In this work, we investigate the

The work was supported by EPSRC under grants EP/T028270/1 and ST/W000970/1. Computing resources came from the Cirrus UK National Tier-2 HPC Service at EPCC (<http://www.cirrus.ac.uk>) funded by the University of Edinburgh and EPSRC (EP/P020267/1).

¹<https://public.nrao.edu/telescopes/vla/>

robustness of the R2D2 image estimation process, studying the uncertainty associated with its series of learned models.

The remainder of this paper is structured as follows. In Section II, the methodology underlying the joint image and estimation and uncertainty quantification with R2D2 is summarized. Section III provides a the validation of the proposed methodology, focusing on RI imaging, and adopting a telescope-specific approach. We conclude in Section IV

II. R2D2 METHODOLOGY

A. Image Reconstruction

At the core of R2D2 lies a series of DNNs ($\mathbf{N}_{\hat{\theta}^{(i)}})_{1 \leq i \leq I}$, each wielding the power of deep learning to enhance image estimates through learned parameters $\hat{\theta}^{(i)} \in \mathbb{R}^Q$. At each iteration $i \in \{1, \dots, I\}$, the network $\mathbf{N}_{\hat{\theta}^{(i)}}$ takes as input both the previous image estimate $\mathbf{x}^{(i-1)}$ and its associated residual dirty image, given by:

$$\mathbf{r}^{(i-1)} = \mathbf{x}_d - \kappa \text{Re}\{\Phi^\dagger \Phi\} \mathbf{x}^{(i-1)}. \quad (3)$$

At the first iteration, the initial image estimate and residual dirty image are initialized as $\mathbf{x}^{(0)} = 0 \in \mathbb{R}^N$ and $\mathbf{r}^{(0)} = \mathbf{x}_d$. Subsequently, the image estimate is updated by adding the output of the network to the previous estimate:

$$\mathbf{x}^{(i)} = \mathbf{x}^{(i-1)} + \mathbf{N}_{\hat{\theta}^{(i)}}(\mathbf{r}^{(i-1)}, \mathbf{x}^{(i-1)}). \quad (4)$$

This iterative refinement gradually improves the resolution and dynamic range of the reconstruction. Consequently, the final reconstructed image $\hat{\mathbf{x}}$ is a “series” formed by the summation of output residual images from all network components.

Training a sequence of R2D2 DNNs ($\mathbf{N}_{\hat{\theta}^{(i)}})_{1 \leq i \leq I}$ commences with a dataset comprising K ground truth images \mathbf{x}_k^* , from which simulated measurements ($\mathbf{y}_k)_{1 \leq k \leq K}$ are generated following the procedure outlined in (1). Subsequently, at each $1 \leq i \leq I$, a set of residuals $\mathbf{r}_k^{(i-1)}$ is created via (3). Each DNN is trained by minimizing the loss function defined as:

$$\min_{\theta^{(i)} \in \mathbb{R}^Q} \frac{1}{K} \sum_{k=1}^K \|\mathbf{x}_k^* - [\mathbf{x}_k^{(i-1)} + \mathbf{N}_{\theta^{(i)}}(\mathbf{r}_k^{(i-1)}, \mathbf{x}_k^{(i-1)})]_+\|_1, \quad (5)$$

where $\|\cdot\|_1$ denotes the ℓ_1 -norm. To maintain the positivity of the image estimate during each iteration, the operator $[\cdot]_+$ projects the image estimate onto the positive orthant \mathbb{R}_+^N .

B. Robustness & Uncertainty Quantification

Uncertainty quantification is of paramount importance in ill-posed inverse imaging problems. On the one hand, the incompleteness of the observed data unavoidably leads to uncertainty, aka aleatoric uncertainty. It can typically be studied by sampling from the posterior distribution $p(\mathbf{x}|\mathbf{y})$, derived from (1) using Bayes’ theorem, with a sampling algorithm, such as Monte Carlo Markov Chain (MCMC) methods [15], [16]. The uncertainty can then be described by the statistics of the samples, such as their mean and standard deviation. On the other hand, the uncertainty associated with the choice of a specific regularization model, aka epistemic uncertainty, is no less important.

The R2D2 algorithm is deterministic in nature, in which context the quantification of aleatoric uncertainty is not directly accessible. We instead propose in this work to investigate the epistemic uncertainty in the reconstruction, associated with the choice regularization model encapsulated in the DNNs. Adopting an ensemble averaging approach, multiple series can be trained, arising from different random DNN initializations of the training process at each iteration. The resulting multiple R2D2 instances can also be leveraged to generate “R2D2 samples”, from which empirical mean and standard deviation endow the algorithm with a joint estimation and uncertainty quantification functionality. Similar epistemic uncertainty quantification approaches have been proposed in [17]–[19].

III. VALIDATION

A. Training Dataset and Training

As in [14], our ground truth dataset comprises 20000 images of size $N = 512 \times 512$ and of peak pixel value 1, curated from astronomical images at optical wavelengths, as well as medical images, pre-processed following [17]. We define the dynamic range a of an image of interest as the ratio of its peak to the faintest pixel intensity, or similarly the standard deviation of the background noise. In order to ensure that the range of dynamic ranges present in the ground truth dataset matches the extremely high values expected in RI, typically $a \in [10^3, 5 \times 10^5]$, pixel-wise exponentiation procedures are applied, as described in [10]. Adopting the same VLA-specific approach undertaken in [14], we sample a different Fourier sampling pattern per each ground truth image, resulting in 20000 dirty images as per (2). The measurement noise is adapted to ensure that the back-projected noise level is commensurate to the inverse dynamic range of the exponentiated ground truth image. A comprehensive depiction of the training procedure, including data fidelity-driven pruning of the training dataset, is outlined in [14]. Our core network architecture for each DNN is the popular U-Net DNN [20], optimized for each R2D2 iteration using (5). Specific to the target uncertainty quantification analysis, and targeting a simple proof of concept, we train $L = 5$ independent realizations of the R2D2 DNN series.

B. Simulations: Test Dataset

The test dataset includes four authentic radio images, namely 3C353 (sourced from NRAO Archives), Messier 106 [21], Abell 2034 and PSZ2 G165.68+44.01 [22], from where 200 ground truth images of size $N = 512 \times 512$ were generated following the same pre-processing and exponentiation steps employed for the training dataset. Corresponding dirty images \mathbf{x}_d were created by applying an unique sampling pattern to each ground truth image via (2).

C. Simulations: Benchmark algorithms

The joint R2D2 functionality for image estimation and uncertainty quantification will be validated, in comparison to: (i) a U-Net trained end-to-end to deliver a reconstruction directly

from the dirty image (which is nothing else than R2D2's first iteration); (ii) uSARA, AIRI, and CLEAN (delivering image estimation only).

D. Simulations: Metrics

We assess image reconstruction quality using the signal-to-noise ratio (SNR) metric, given by:

$$\text{SNR}(\hat{\mathbf{x}}, \mathbf{x}^*) = 20\log_{10}(\|\mathbf{x}^*\|_2 / \|\mathbf{x}^* - \hat{\mathbf{x}}\|_2), \quad (6)$$

where $\hat{\mathbf{x}}$ is the reconstruction by any algorithm with ground truth \mathbf{x}^* . To examine the performance of an algorithm in high-dynamic range scenario and to facilitate visualization of faint features, we transform an image \mathbf{x} to logarithmic scale by:

$$\text{rlog}(\mathbf{x}) = x_{\max} \log_a \left(\frac{a}{x_{\max}} \mathbf{x} + \mathbf{1} \right), \quad (7)$$

where a is the dynamic range of the image of interest and x_{\max} is the peak pixel value of \mathbf{x} . The SNR metric in logarithmic scale is given by:

$$\log\text{SNR}(\hat{\mathbf{x}}, \mathbf{x}^*) = \text{SNR}(\text{rlog}(\hat{\mathbf{x}}), \text{rlog}(\mathbf{x}^*)). \quad (8)$$

The set of image reconstructions from the $L = 5$ realizations of R2D2 read:

$$\bar{\mathbf{x}}^{(i)} = [\hat{\mathbf{x}}_1^{(i)}, \dots, \hat{\mathbf{x}}_L^{(i)}], \quad (9)$$

where $\hat{\mathbf{x}}_\ell^{(i)}$ is the estimate from the ℓ -th R2D2 realization at the i -th iteration. To investigate the associated uncertainty, we evaluate the pixel-wise mean $\mu(\bar{\mathbf{x}})$, given by:

$$\mu(\bar{\mathbf{x}}^{(i)}) = \frac{1}{L} \sum_{\ell=1}^L \hat{\mathbf{x}}_\ell^{(i)}, \quad (10)$$

as well as the pixel-wise ratio of the standard deviation to mean, denoted as $[\sigma/\mu](\bar{\mathbf{x}})$, which reads:

$$[\sigma/\mu](\bar{\mathbf{x}}^{(i)}) = \begin{cases} \frac{1}{\mu(\bar{\mathbf{x}}^{(i)})} \sqrt{\frac{\sum_{\ell=1}^L (\hat{\mathbf{x}}_\ell^{(i)} - \mu(\bar{\mathbf{x}}^{(i)}))^2}{L}} & \text{if } \mu(\bar{\mathbf{x}}^{(i)}) < 1/a, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

E. Simulations: Results & Discussion

Fig. 1 illustrates the progression of the R2D2 series from its initial to final iteration, simultaneously peeling residuals $\mathbf{r}_{\text{R2D2}}^{(i)}$ and updating the image estimate $\hat{\mathbf{x}}_{\text{R2D2}}^{(i)}$ from the DNN output.

The left and center panels of Fig. 2 present the graphs of the SNR and logSNR metrics across iterations, respectively. For R2D2, these metrics are computed both for the image estimate $\hat{\mathbf{x}}_{\text{R2D2}}^{(i)}$ of a single R2D2 realization and for their pixel-wise mean $\mu(\bar{\mathbf{x}}^{(i)})_{\text{R2D2}}$. The U-Net results are those of the first R2D2 iteration. Final results of uSARA and CLEAN, respectively denoted $\hat{\mathbf{x}}_{\text{uSARA}}$ and $\hat{\mathbf{x}}_{\text{CLEAN}}$, are shown as horizontal lines. For simplicity, AIRI results, on par with those of uSARA, are not reported. The mean and standard deviation across the 200 test inverse problems and 5 R2D2 realizations are reported for $\mu(\bar{\mathbf{x}}^{(i)})_{\text{R2D2}}$, and simply across the 200 test inverse problems for $\hat{\mathbf{x}}_{\text{R2D2}}^{(i)}$, $\hat{\mathbf{x}}_{\text{uSARA}}$, and $\hat{\mathbf{x}}_{\text{CLEAN}}$. Firstly, the results confirm that R2D2's image estimation capability is




Fig. 1: Evolution of the R2D2 series (single realization) across its iterations. The ground truth image \mathbf{x}^* and dirty image $\mathbf{r}_{\text{R2D2}}^{(0)}$ are reported in the top left and center panels, respectively. The output of the first DNN gives $\hat{\mathbf{x}}_{\text{R2D2}}^{(1)}$ (top right). 12 iterations later, $\mathbf{r}_{\text{R2D2}}^{(11)}$ (bottom left) and $\hat{\mathbf{x}}_{\text{R2D2}}^{(12)}$ are fed to the last DNN, whose output (bottom center) is added to $\hat{\mathbf{x}}_{\text{R2D2}}^{(12)}$ to deliver the reconstruction (bottom right).

superior to that of the state-of-the-art algorithms. Secondly, in both SNR and logSNR, $\mu(\bar{\mathbf{x}}^{(i)})_{\text{R2D2}}$ yields slightly superior results compared to $\hat{\mathbf{x}}_{\text{R2D2}}^{(i)}$, suggesting an advantage in averaging “R2D2 samples”, in terms of precision.

Table I reports the computational cost for different imaging algorithms, in terms of number of iterations required and total reconstruction time. In addition to the higher SNR and logSNR values shown in Fig. 2 when compared to the benchmarks, one can note that R2D2 also offers significant computational efficiency, requiring drastically fewer iterations than uSARA and AIRI, and taking only a fraction of time required by CLEAN, thanks to its ultra-fast DNN-encapsulated regularization. We note that U-Net and R2D2 are executed on a single GPU, CLEAN and uSARA are executed on a single CPU, while AIRI is implemented with regularization DNNs on GPU and data-fidelity steps on CPU. These results confirm that R2D2 boasts ultra-fast reconstruction capability, which makes the computation of “R2D2 samples” practical even at large image dimension.

The graph of the relative pixel-wise epistemic uncertainty $[\sigma/\mu](\bar{\mathbf{x}}^{(i)})_{\text{R2D2}}$, more specifically its mean and standard deviation across the 200 test inverse problems, is presented in the right panel of Fig. 2. A consistent decrease of $[\sigma/\mu](\bar{\mathbf{x}}^{(i)})_{\text{R2D2}}$ across the R2D2 iterations is observed. This observation serves as compelling evidence that the R2D2 robustness (resp. epistemic uncertainty) increases (resp. decreases) rapidly with the number of terms in its series.

In Fig. 3, estimated images are presented for the various algorithms under scrutiny, and associated relative uncertainty




Fig. 2: Quantitative metric analysis. Left (resp. center): Progress of the reconstruction quality evaluated by SNR (resp. logSNR) across R2D2 iterations for both $\mu(\tilde{\mathbf{x}}^{(i)})_{\text{R2D2}}$ and $\hat{\mathbf{x}}_{\text{R2D2}}^{(i)}$. Right: Progress of the epistemic uncertainty evaluated by $[\sigma/\mu](\tilde{\mathbf{x}}^{(i)})_{\text{R2D2}}$ across iterations. Values of the SNR and logSNR, achieved at convergence by the benchmark algorithms uSARA and CLEAN, are reported via horizontal lines, in their respective plots. Each point is the averaged metric over 200 inverse problems (5×200 for $\hat{\mathbf{x}}_{\text{R2D2}}$), while the shaded area presents its standard deviations.

Algorithm	$I \pm \text{std}$	$t_{\text{tot.}} \pm \text{std (sec)}$
$\hat{\mathbf{x}}_{\text{U-Net}}$	1	0.1 ± 0.1
$\hat{\mathbf{x}}_{\text{R2D2}}$	12	2.0 ± 0.4
$\hat{\mathbf{x}}_{\text{AIRI}}$	4995 ± 50	3430 ± 1461
$\hat{\mathbf{x}}_{\text{uSARA}}$	1107 ± 377	4015.9 ± 1471.2
$\hat{\mathbf{x}}_{\text{CLEAN}}$	8 ± 1	93.2 ± 26.7

TABLE I: Computational cost of each imaging algorithm used in this study, in terms of number of iterations I and total reconstruction time $t_{\text{tot.}}$ in seconds.

maps $[\sigma/\mu](\tilde{\mathbf{x}}^{(i)})_{\text{R2D2}}$ for the R2D2 model, for one of the 200 inverse problems in the test dataset. The first row displays the ground truth \mathbf{x}^* (left) and the dirty image \mathbf{x}_d (right), the second row $\hat{\mathbf{x}}_{\text{uSARA}}$ (left), and $\hat{\mathbf{x}}_{\text{CLEAN}}$ (right). The third (resp. forth, fifth) row displays $\mu(\tilde{\mathbf{x}}^{(i)})_{\text{R2D2}}$ (left) and $[\sigma/\mu](\tilde{\mathbf{x}}^{(i)})_{\text{R2D2}}$ (right) at iterations $i = 1$ (U-Net) (resp. $i = 3$, $i = 12$). AIRI results, on par with those of uSARA, are not displayed. For the R2D2 uncertainty maps, the mean pixel value is reported on the image itself, highlighting a relative uncertainty of the order 0.2% only for R2D2, against 5% for U-Net. These results corroborate the analysis of the metrics in Fig. 2.

F. Real Data Proof of Concept

In this section, the joint image estimation and uncertainty quantification functionality of R2D2 is validated on real data. R2D2 is applied to image of the radio galaxy Cygnus A from S band observations with the Very Large Array (VLA; see [13] for detailed data description). The left panel of Fig. 4 displays the pixel-wise mean image $\mu(\tilde{\mathbf{x}}^{(12)})_{\text{R2D2}}$ (zoom on the right lobe of the galaxy) derived from 5 reconstructions generated by the R2D2 algorithm. The reconstruction exhibits high resolution across both low and high intensity values, illustrating the algorithm’s precision. The right panel displays the associated uncertainty map $[\sigma/\mu](\tilde{\mathbf{x}}^{(12)})_{\text{R2D2}}$ (zoom on the

same region). The mean pixel value is reported in the upper left corner, highlighting a relative uncertainty of the order 0.5% only.

IV. CONCLUSION

In this work, we have investigated the robustness of the R2D2 image estimation process, adopting an ensemble averaging approach and studying the uncertainty arising from the random initialization of the training process for each DNN of the series. Focusing on high-dynamic range RI imaging in astronomy, 5 R2D2 series realizations with 12 terms each were trained for the VLA. Simulations and real-data experiments confirm that: (i) R2D2’s image estimation capability is largely superior to that of CLEAN, U-Net, and even uSARA and AIRI; (ii) R2D2 boasts ultra-fast reconstruction capability, making the computation of multiple reconstruction samples and of uncertainty maps practical even at large image dimension; (iii) it is characterized by a very low epistemic uncertainty (contrasting to the first iteration of the series, *i.e.* U-Net).

REFERENCES

- [1] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy*. Springer Nature, 2017.
- [2] D. Briggs, “High fidelity interferometric imaging: Robust weighting and nnls deconvolution,” in *AAS*, vol. 187, 1995, pp. 112–02.
- [3] J. Högbom, “Aperture synthesis with a non-regular distribution of interferometer baselines,” *A&AS*, vol. 15, p. 417, 1974.
- [4] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [5] Y. Wiaux, L. Jacques, G. Puy, A. M. Scaife, and P. Vanderghelynst, “Compressed sensing imaging techniques for radio interferometry,” *MNRAS*, vol. 395, no. 3, pp. 1733–1742, 2009.
- [6] F. Li, T. J. Cornwell, and F. de Hoog, “The application of compressive sampling to radio astronomy-I. Deconvolution,” *A&A*, vol. 528, p. A31, 2011.
- [7] R. E. Carrillo, J. D. McEwen, and Y. Wiaux, “Sparsity Averaging Reweighted Analysis (SARA): a novel algorithm for radio-interferometric imaging,” *MNRAS*, vol. 426, no. 2, pp. 1223–1234, 2012.

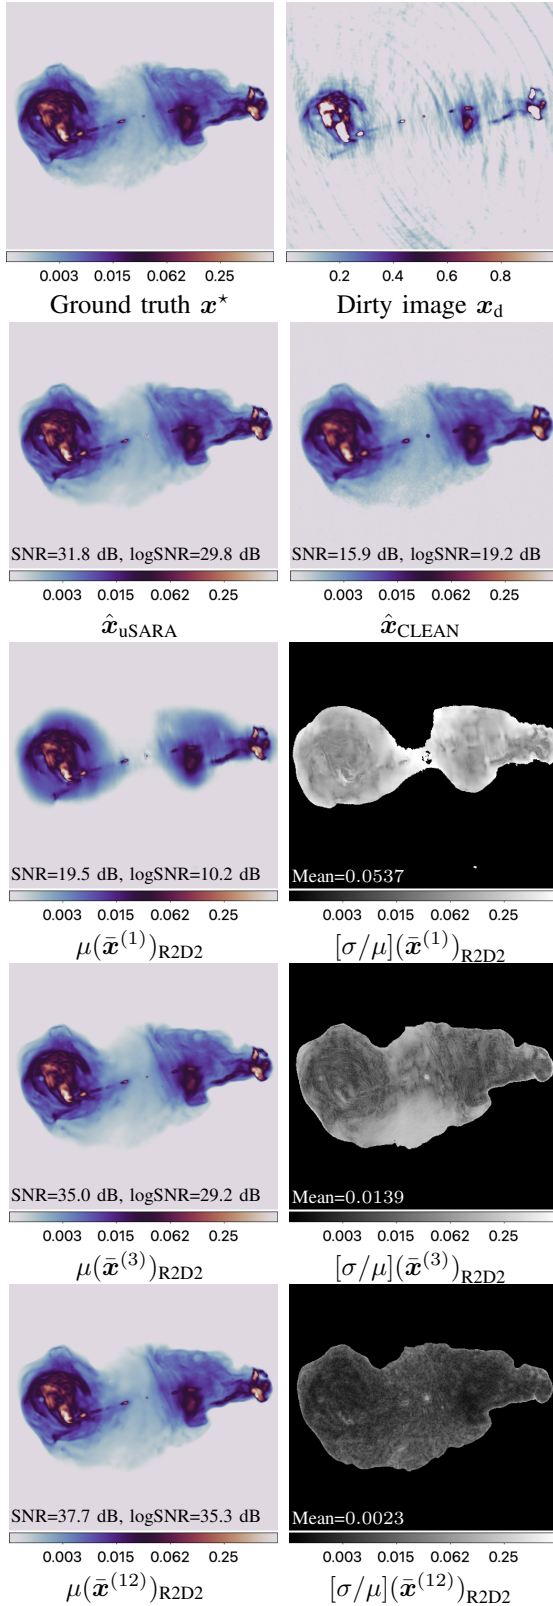


Fig. 3: Visual illustration of the R2D2 joint image estimation and uncertainty quantification functionality across iterations for the simulation results. Top row: ground truth (left) and dirty image (right). Second row: reconstruction \hat{x}_{uSARA} (left) and \hat{x}_{CLEAN} (right). Third (resp. fourth, fifth) row: $\mu(\bar{x}^{(i)})_{R2D2}$ (left) and $[\sigma/\mu](\bar{x}^{(i)})_{R2D2}$ (right) for iterations $i = 1$ (U-Net) (resp. $i = 3$, $i = 12$). The evaluation metrics (SNR, logSNR) are embedded inside each image estimate. The mean pixel values across the image is shown inside each uncertainty image.




Fig. 4: Visual illustration of the R2D2 joint image estimation and uncertainty quantification functionality across iterations for the real data results. Left: pixel-wise mean $\mu(\bar{x}^{(12)})_{R2D2}$. Right: uncertainty map $[\sigma/\mu](\bar{x}^{(12)})_{R2D2}$.

- [8] A. Dabbech, C. Ferrari, D. Mary, E. Slezak, O. Smirnov, and J. S. Kenyon, "Moresane: Model REconstruction by Synthesis-ANalysis Estimators-a sparse deconvolution algorithm for radio interferometric imaging," *A&A*, vol. 576, p. A7, 2015.
- [9] A. Repetti and Y. Wiaux, "A forward-backward algorithm for reweighted procedures: application to radio-astronomical imaging," in *Proc. IEEE ICASSP 2020*, 2020, pp. 1434–1438.
- [10] M. Terris, A. Dabbech, C. Tang, and Y. Wiaux, "Image reconstruction algorithms in radio interferometry: From handcrafted to learned regularization denoisers," *MNRAS*, vol. 518, no. 1, pp. 604–622, 09 2022.
- [11] L. Connor, K. L. Bouman, V. Ravi, and G. Hallinan, "Deep radio-interferometric imaging with POLISH: DSA-2000 and weak lensing," *MNRAS*, vol. 514, no. 2, pp. 2614–2626, 2022.
- [12] Schmidt, K., Geyer, F., Fröse, S., Blomenkamp, P.-S., Brüggem, M., de Gasperin, F., Elsässer, D., and Rhode, W., "Deep learning-based imaging in radio interferometry," *A&A*, vol. 664, p. A134, 2022.
- [13] A. Aghabiglou, C. S. Chu, A. Jackson, A. Dabbech, and Y. Wiaux, "Ultra-fast high-dynamic range imaging of cygnus a with the R2D2 deep neural network series," *preprint arXiv:2309.03291*, 12 2023.
- [14] A. Aghabiglou, S. C. Chu, A. Dabbech, and Y. Wiaux, "The R2D2 deep neural network series paradigm for fast precision imaging in radio astronomy," *preprint arXiv:2403.05452*, 2024.
- [15] R. Laumont, V. D. Bortoli, A. Almansa, J. Delon, A. Durmus, and M. Pereyra, "Bayesian imaging using plug & play priors: when langevin meets tweedie," *SIAM J. Imaging Sci.*, vol. 15, no. 2, pp. 701–737, 2022.
- [16] S. Mukherjee, A. Hauptmann, O. Öktem, M. Pereyra, and C.-B. Schönlieb, "Learned reconstruction methods with convergence guarantees: A survey of concepts and applications," *IEEE Signal Process. Mag.*, vol. 40, no. 1, pp. 164–182, 2023.
- [17] M. Terris, C. Tang, A. Jackson, and Y. Wiaux, "Plug-and-play imaging with model uncertainty quantification in radio astronomy," *preprint arXiv:2312.07137*, 2023.
- [18] D. Narnhofer, A. Effland, E. Kobler, K. Hammernik, F. Knoll, and T. Pock, "Bayesian uncertainty estimation of learned variational mri reconstruction," *IEEE Trans. Med. Imaging*, vol. 41, no. 2, pp. 279–291, 2021.
- [19] S. Lahlou, M. Jain, H. Nekoei, V. I. Butoi, P. Bertin, J. Rector-Brooks, M. Korablyov, and Y. Bengio, "Deup: Direct epistemic uncertainty prediction," *preprint arXiv:2102.08501*, 2021.
- [20] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *MICCAI 2015*, 2015, pp. 234–241.
- [21] T. Shimwell, M. Hardcastle, C. Tasse, P. Best, H. Röttgering, W. Williams, A. Botteon, A. Drabent, A. Mechev, A. Shulevski *et al.*, "The LOFAR Two-metre Sky Survey-V. Second data release," *A&A*, vol. 659, p. A1, 2022.
- [22] A. Botteon, T. Shimwell, R. Cassano, V. Cuciti, X. Zhang, L. Bruno, L. Camillini, R. Natale, A. Jones, F. Gastaldello *et al.*, "The Planck clusters in the LOFAR sky-I. LoTSS-DR2: New detections and sample overview," *A&A*, vol. 660, p. A78, 2022.