

R2D2 FOR LARGE-SCALE RADIO-INTERFEROMETRIC IMAGING

YVES WIAUX

on behalf of BASP & Co.

SEPTEMBER, 2026

SWISS SKA DAYS, NEUCHÂTEL, SWITZERLAND

☞ MODERN CHALLENGES IN RADIO-ASTRONOMICAL IMAGING

↪ *On the need to go beyond CLEAN*

☞ THE R2D2 CONCEPT

↪ *A learned version of CLEAN*

☞ TOWARDS SCALABLE R2D2

↪ *Enabling wide-field imaging*

MODERN CHALLENGES IN RADIO-ASTRONOMICAL IMAGING

↪ *On the need to go beyond CLEAN*

Across science and technology, when observation gathers incomplete information about an image, **advanced computational imaging algorithms are needed to transform data into images.**

- ▶ **Applications** are ubiquitous, but it all quite started with deconvolution problems in astronomy...

Challenges:

- ▶ Precision: high **resolution** and **dynamic range**
- ▶ Robustness: **uncertainty quantification** and **calibration**
- ▶ Efficiency: **fast acquisition** and **reconstruction**
- ▶ Scalability: **large data volumes**

The Square Kilometre Array

SKA will target unprecedented resolution and sensitivity regimes, leading to **EB data volumes** and **PB wide-band image sizes**.





Image credit SKA organisation

- Reconstruction algorithms must be **efficient** and **scalable**.

Aperture synthesis by radio interferometry

Aperture synthesis in radio interferometry (RI) maps the radio sky at high resolutions and sensitivities through **spatial Fourier sampling**.

- The incomplete Fourier sampling leads to an **ill-posed linear inverse problem** to be solved:




The dirty image x_{dirty} can be thought as a convolution of x^* with PSF $\Phi^\dagger \Phi \delta$:

$$x_{\text{dirty}} = \Phi^\dagger y \simeq \Phi^\dagger \Phi \delta \star x^* + n'$$

Aperture synthesis by radio interferometry

Solving the **inverse problem**, i.e, recovering $\mathbf{x}^*$ from $\mathbf{y}$ can be thought as solving an **image deconvolution problem**, i.e, recovering $\mathbf{x}^*$ from $\mathbf{x}_{\text{dirty}}$.



- **Reconstruction algorithms** are needed, leveraging a **prior** image model to **regularise** and solve the inverse problem.

The standard **CLEAN** algorithm is a **greedy matching pursuit algorithm**, iteratively identifying model components from back-projected data residuals.

► **CLEAN iterative scheme:**

- Initialisation:

$$\mathbf{x}^{(0)} = \mathbf{0}; \mathbf{r}^{(0)} = \mathbf{x}_{\text{dirty}}$$

- For $i = 0, \dots, I$: **(major cycle)**

$$\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \boxed{\mathcal{T}(\mathbf{r}^{(i)})} \longrightarrow \text{minor cycle}$$

$$\mathbf{r}^{(i+1)} = \mathbf{x}_{\text{dirty}} - \Phi^\dagger \Phi \mathbf{x}^{(i+1)}.$$

- $\mathcal{T}(\cdot)$ **peeling operator** implicitly enforcing a **sparse image model**.

The standard **CLEAN** algorithm is a **greedy matching pursuit algorithm**, iteratively identifying model components from back-projected data residuals.

- ▶ CLEAN iterative scheme: $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \mathcal{T}(\mathbf{r}^{(i)})$
- ▶ $\mathcal{T}(\cdot)$ **peeling operator** implicitly enforcing a **sparse image model**.
- ▶ **Simplistic model**: bringing efficiency and scalability, but limiting precision and robustness.
- ▶ **RI image reconstruction is to be reinvented**.

THE R2D2 CONCEPT

↪ *A learned version of CLEAN*

REFERENCES

↔ AGHABIGLOU ET AL., PROC. ICASSP 2023, ARXIV:2210.1606

↔ DABBECH ET AL., APJL, 2024, ARXIV:2309.03291

↔ AGHABIGLOU ET AL., APJS, 2024, ARXIV:2403.05452

↔ AGHABIGLOU ET AL., PROC. EUSIPCO, ARXIV:2403.18052

↔ AGHABIGLOU ET AL., APJS, 2025, ARXIV:2503.02554



A learned version of CLEAN

Residual-to-Residual DNN series for high Dynamic-range imaging: R2D2 applies DNNs iteratively, each inputting the current back-projected data residual and outputting an image residual... progressively peeling the dynamic range.

► **R2D2 iterative scheme:**

- Initialisation:

$$\mathbf{x}^{(0)} = \mathbf{0}; \mathbf{r}^{(0)} = \mathbf{x}_{\text{dirty}}$$

- For $i = 0, \dots, \mathcal{I}$: (**major cycle**)

$$\mathbf{x}^{(i+1)} = \left[\mathbf{x}^{(i)} + \mathbf{N}_{\hat{\boldsymbol{\theta}}^{(i+1)}}(\mathbf{r}^{(i)}, \mathbf{x}^{(i)}) \right]_+ \longrightarrow \text{minor cycle}$$

$$\mathbf{r}^{(i+1)} = \mathbf{x}_{\text{dirty}} - \boldsymbol{\Phi}^\dagger \boldsymbol{\Phi} \mathbf{x}^{(i+1)}.$$

- $\mathbf{N}_{\hat{\boldsymbol{\theta}}^{(i+1)}}$ is a U-Net-type convolutional neural network.
- $\mathbf{N}_{\hat{\boldsymbol{\theta}}^{(i+1)}}$ is iteration dependent.

A learned version of CLEAN

Residual-to-Residual DNN series for high Dynamic-range imaging: R2D2 applies DNNs iteratively, each inputting the current back-projected data residual and outputting an image residual... progressively peeling the dynamic range.

► **R2D2 iteration structure:**

$$\mathbf{x}^{(i+1)} = [\mathbf{x}^{(i)} + \mathbf{N}_{\hat{\theta}^{(i+1)}}(\mathbf{r}^{(i)}, \mathbf{x}^{(i)})]_+$$


with $\mathbf{x}^{(0)} = \mathbf{0}$, $\mathbf{r}^{(i)} = \mathbf{x}_{\text{dirty}} - \kappa \text{Re}\{\mathbf{\Phi}^\dagger \mathbf{\Phi}\} \mathbf{x}^{(i)}$.

- **Supervised training** of $\mathbf{N}_{\hat{\theta}^{(i+1)}}$ at iteration $i + 1$, using the reconstruction and data-residual training datasets of iteration i , $(\mathbf{x}_k^{(i)}, \mathbf{r}_k^{(i)})_k$:

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

R2D2 DNNs are trained on **synthetic high-dynamic range databases**, with flexibility in the underpinning network architecture.


► **U-WDSR** architecture underpinning R2D2 for static imaging:



DNN architecture and training database

R2D2 DNNs are trained on **synthetic high-dynamic range databases, with flexibility in the underpinning network architecture.**


- Creating a high-dynamic range database by exponentiating low-dynamic range astronomical and medical image datasets:



Low-dynamic range database

- ✓ artefacts removal
- ✓ size adjustment 512×512
- ✓ **exponentiated** for dynamic range in $[10^3, 5 \times 10^5]$

Pre-processing



High-dynamic range database





Training R2D2 for VLA

A first instance of R2D2 is **specifically trained and validated for VLA imaging**.

► **Training vs. validation images:** $(\mathbf{x}_k^*)_k$

Training samples

Validation samples (4 **real radio** images)



Astronomical images

Medical images

3c353

M106

PSZ2

ACO 2034

► **Observation model:** $(\Phi_k)_k$

- ✓ 20k VLA sampling patterns with [0.2, 27] million data points
- ✓ Briggs weighting $[-1, 1]$
- ✓ Super-resolution factor $[1.5, 2.5]$
- ✓ Training and validation dirty images generated using $\mathbf{x}_k^{\text{dirty}} = \Phi_k^\dagger(\Phi_k \mathbf{x}_k^* + \mathbf{n}_k)$

Visual simulation results


Reconstruction results from simulated observations of the 3C353 source.

► **Model evolution** across R2D2 iterations:

$\mathbf{r}^{(0)}$


$\mathbf{N}_{\hat{\theta}^{(1)}}(\mathbf{r}^{(0)}, \mathbf{x}^{(0)})$

$\mathbf{x}^{(1)}$




0.048 1.2 2.3 3.3 4.5

Linear scale visualisation



0.00099 0.015 0.12

Log scale visualisation



0.00099 0.015 0.12

Log scale visualisation

Visual simulation results


Reconstruction results from simulated observations of the 3C353 source.

► **Model evolution** across R2D2 iterations:

$r^{(1)}$


$N_{\hat{\theta}^{(2)}}(r^{(1)}, x^{(1)})$

$x^{(2)}$




-0.4 -0.1 0.2

Linear scale visualisation



-0.032 -0.016 3.9e-05 0.016 0.032

Linear scale visualisation



0.00099 0.015 0.12

Log scale visualisation

Visual simulation results


Reconstruction results from simulated observations of the 3C353 source.

► **Model evolution** across R2D2 iterations:

$\mathbf{r}^{(13)}$


$\mathbf{N}_{\hat{\theta}^{(14)}}(\mathbf{r}^{(13)}, \mathbf{x}^{(13)})$

$\mathbf{x}^{(14)}$




-0.0048 -0.0012 0.0024

Linear scale visualisation



-0.0008 -0.0002 0.0004

Linear scale visualisation



0.00099 0.015 0.12

Log scale visualisation

Quantitative metrics confirm **R2D2** brings superior precision to optimisation and plug-and-play algorithms... at a fraction of the cost (efficiency).

► Simulation results:

Model	Image quality metrics		DATA fidelity	# Iterations	Run time
	SNR $\pm$ std (dB)	logSNR $\pm$ std (dB)	RDR* $\pm$ std ($\times 10^{-3}$)	I $\pm$ std	$t_{\text{tot.}}$ $\pm$ std (s)
CLEAN	12.0 $\pm$ 19.3	9.4 $\pm$ 19.0	3.29 $\pm$ 2.8	8.4 $\pm$ 1.0	106.5 $\pm$ 81.6
uSARA	28.1 $\pm$ 3.4	20.4 $\pm$ 3.4	2.15 $\pm$ 2.8	1482 $\pm$ 586	368.8 $\pm$ 296.8
AIRI	28.3 $\pm$ 3.1	21.1 $\pm$ 3.8	2.24 $\pm$ 2.8	5000 $\pm$ 0.0	937.4 $\pm$ 801.8
U-WDSR	16.0 $\pm$ 3.6	6.6 $\pm$ 3.8	155.8 $\pm$ 81.6	1	0.66 $\pm$ 0.03
R2D2	31.3 $\pm$ 2.4	24.6 $\pm$ 4.16	2.22 $\pm$ 28.1	15.8 $\pm$ 5.5	8.8 $\pm$ 3.9

* : Residual-to-Dirty image Ratio

TOWARDS SCALABLE R2D2

↪ *Enabling wide-field imaging*

REFERENCES

↔ AGHABIGLOU ET AL., APJS IN PREPARATION, 2026



Residual-to-Residual DNN series for high Dynamic-range imaging: R2D2 applies DNNs iteratively, each inputting the current back-projected data residual and outputting an image residual... progressively peeling the dynamic range.

► **R2D2 iteration structure:**

$$\mathbf{x}^{(i+1)} = [\mathbf{x}^{(i)} + \mathbf{N}_{\hat{\theta}^{(i+1)}}(\mathbf{r}^{(i)}, \mathbf{x}^{(i)})]_+$$

with $\mathbf{x}^{(0)} = \mathbf{0}$, $\mathbf{r}^{(i)} = \mathbf{x}_{\text{dirty}} - \kappa \text{Re}\{\mathbf{\Phi}^\dagger \mathbf{\Phi}\} \mathbf{x}^{(i)}$.

► **Supervised training** of $\mathbf{N}_{\hat{\theta}^{(i+1)}}$ at iteration $i + 1$, using the reconstruction and data-residual training datasets of iteration i , $(\mathbf{x}_k^{(i)}, \mathbf{r}_k^{(i)})_k$:


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

R2D2's novel DNN architecture

R2D2 seamlessly generalises from moderate to **extreme image size** via a **transformer** architecture and **preconditioning scheme**.

- **Transformer-based DNN** architecture underpinning R2D2 for static imaging:

Multi-Briggs input




- Using a **preconditioning scheme** R2D2 marginalising over **Briggs weighting** variables to jointly optimise the **resolution** and **sensitivity**.

MegaSky–Extreme-scale Radio Interferometric Image Generation Pipeline.

► Initialize & populate the target field:


- ✓ create an empty canvas at the desired resolution
- ✓ randomly **transform** radio sources through **scaling, stretching, rotation, local intensity transformation, and tapering**
- ✓ randomly place and add the transformed sources to form a large-scale radio sky model



MegaSky–Extreme-scale Radio Interferometric Image Generation Pipeline.

► Initialize & populate the target field:

- ✓ create an empty canvas at the desired resolution
- ✓ randomly **transform** radio sources through **scaling, stretching, rotation, local intensity transformation, and tapering**
- ✓ randomly place and add the transformed sources to form a large-scale radio sky model




Extreme-scale dataset generation pipeline

MegaSky–Extreme-scale Radio Interferometric Image Generation Pipeline.

► **Initialize & populate the target field:**

- ✓ create an empty canvas at the desired resolution
- ✓ randomly **transform** radio sources through **scaling, stretching, rotation, local intensity transformation, and tapering**
- ✓ randomly place and add the transformed sources to form a large-scale radio sky model




Training R2D2 for VLA at scale


R2D2 learns across **multiple Briggs weightings** to marginalize over imaging choices.

► Training vs. validation images:


Training samples



1024 × 1024




2048 × 2048




4096 × 4096


Validation samples from **LOFAR Surveys**



1024 × 1024



2048 × 2048



4096 × 4096


► Observation model:

- ✓ 20k VLA sampling patterns with [0.2, 27] million data points
- ✓ Each observation is normalized by **5 Briggs weighting** $[-1, -0.5, 0, 0.5, 1]$
- ✓ Super-resolution factor $[1.5, 2.5]$

Quantitative simulation results

Quantitative metrics confirm **R2D2** brings superior precision to optimisation and plug-and-play algorithms... at a fraction of the cost (efficiency).

► Simulation results (4096×4096):



RDR: Residual-to-Dirty image Ratio


R2D2 not only outperforms in precision, but also reconstructs images in near **real time at a fraction of the cost**, while remaining largely unaffected by a **64×** increase in image size.

► **Simulation run time (4096×4096):**

Model	Iteration $\pm$ std	$t_{data} \pm \text{std}$ (sec)	$t_{reg.} \pm \text{std}$ (sec)	$t_{tot.} \pm \text{std}$ (sec)
uSARA	2216 $\pm$ 1202	0.11 $\pm$ 0.08	0.64 $\pm$ 0.41	1685.2 $\pm$ 867.9
AIRI	5000 $\pm$ 0.0	0.11 $\pm$ 0.08	0.67 $\pm$ 0.44	3990.0 $\pm$ 427.4
R2D2 (single Briggs)	18	0.35 $\pm$ 0.09	0.33 $\pm$ 0.32	32.36 $\pm$ 9.69
R2D2	12	0.59 $\pm$ 0.33	0.35 $\pm$ 0.18	20.89 $\pm$ 3.76


Cygnus A images formed from **real VLA** observations at **S band**.

► **Dirty** (20MB data, 4096×4096 pixels):




Cygnus A images formed from **real VLA** observations at **S band**.

- **CLEAN** (20MB data, 4096×4096 pixels, $t_{tot.}$: 30 min, Zoomed in):




Cygnus A images formed from **real VLA** observations at **S band**.

- **uSARA** (20MB data, 4096×4096 pixels, $t_{tot.}$: 47 min, Zoomed in):




Cygnus A images formed from **real VLA** observations at **S band**.

► **AIRI** (20MB data, 4096×4096 pixels, $t_{tot.}$: 60 min, **Zoomed in**):



Cygnus A images formed from **real VLA** observations at **S band**.

► **R2D2** (20MB data, 4096×4096 pixels), $t_{tot.}$: 27 sec, Zoomed in:



Conclusion & future work

R2D2 offers a **new regime of joint precision, robustness, efficiency, and scalability** for radio-astronomical image formation.

Upcoming evolutions:

- Endow **R2D2** with the wide-field, hyperspectral, polarization and calibration **functionality**
- Investigate **universal** (vs. telescope-specific) incarnations

Python and Matlab code

- Enjoy BASP's Computational Imaging Library: **BASPLib**
<https://basp-group.github.io/BASPLib/>

Contact

- y.wiaux@hw.ac.uk

