# R2D2's fast precision imaging in radio astronomy (but not only)

Prof. Yves Wiaux, Edinburgh

*on behalf of BASP & Co.*

Swiss SKA Days, September 3rd 2024

# Computational imaging & modern challenges



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

- Applications range from medicine or defence, to astronomy...

Challenges:

- Precision: unprecedented resolution and dynamic range regimes
- Scalability: extreme data volumes
- Robustness: uncertainty quantification, as well as calibration functionalities (not discussed in this talk)

# Outline



☛ COMPUTATIONAL IMAGING IN RADIO ASTRONOMY

*... challenges & CLEAN*

☛ R2D2 ALGORITHM

*... from DNN series to astronomical imaging (but not only)*

# Computational imaging in radio astronomy

# Aperture synthesis by radio interferometry



Aperture synthesis by radio interferometry (RI) provides access to high resolution high-dynamic range. But forming an image $\widehat{\boldsymbol{x}}$ from visibility data $\boldsymbol{y}$ is an ill-posed inverse problem.

- The data provide an incomplete Fourier sampling of the sky, leading to a deconvolution problem:

$$\boxed{\boldsymbol{y} = \boldsymbol{\Phi}\widehat{\boldsymbol{x}} + \boldsymbol{n}}$$

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

$$\boxed{\boldsymbol{y} \rightarrow \widehat{\boldsymbol{x}}}$$

- Accurate image models are needed for **precision** and **scalability**

# The Square Kilometre Array



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

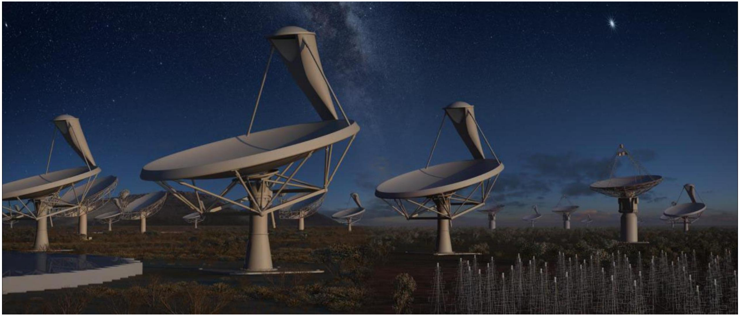

Image credit SKA organisation

- Reconstruction algorithms must be **scalable**

# Celebrated CLEAN

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

- Write backprojected data as convolution of $\widehat{\boldsymbol{x}}$ with PSF $\mathrm{Re}\{\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi}\}\boldsymbol{\delta}$:

$$\boxed{\boldsymbol{x}_{\mathrm{dirty}} = \kappa\mathrm{Re}\{\boldsymbol{\Phi}^\dagger\boldsymbol{y}\} \simeq \kappa\mathrm{Re}\{\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi}\}\boldsymbol{\delta} \star \widehat{\boldsymbol{x}} + \boldsymbol{n}'}$$

with $\kappa = 1/\max(\mathrm{Re}\{\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi}\}\boldsymbol{\delta})$

- CLEAN iteration structure:

$$\boxed{\boldsymbol{x}^{(i)} = \boldsymbol{x}^{(i-1)} + \boldsymbol{\mathcal{T}}\left(\boldsymbol{x}_{\mathrm{dirty}} - \kappa\mathrm{Re}\{\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi}\}\boldsymbol{x}^{(i-1)}\right)}$$

with $\boldsymbol{\mathcal{T}}$ peeling operator implicitly enforcing a sparse image model

- Simplistic model: scalable, but limiting precision
- RI image reconstruction is to be reinvented

# R2D2 algorithm

# 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
- Chen et al., Proc. EUSIPCO, arXiv:2403.17905

# R2D2: a learned version of CLEAN

Non satisfied with the highly iterative nature of optimisation algorithms, hampering scalability, deep learning solutions are appealing.

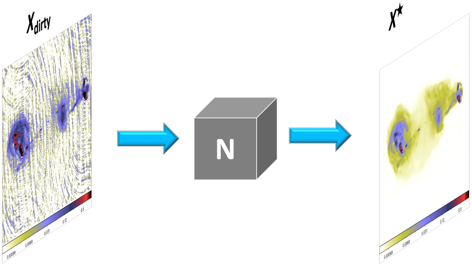


$$\boxed{\boldsymbol{x}^\star = \mathbf{N}(\boldsymbol{x}_{\mathrm{dirty}})}$$

- Purely data-driven DNNs do not generalise well.
- Unfolded DNNs do not scale well due to limitations in embedding measurement operators in network architectures.

# R2D2: a learned version of CLEAN



R2D2 applies DNNs iteratively, each taking the previous iteration's image estimate and back-projected data residual as input, and reconstructing the residual between the ground truth and the reconstruction of the previous iteration.

- R2D2 iteration structure:

$$\boxed{\boldsymbol{x}^{(i)} = \boldsymbol{x}^{(i-1)} + \mathbf{N}_{\widehat{\boldsymbol{\theta}}^{(i)}}(\boldsymbol{r}^{(i-1)}, \boldsymbol{x}^{(i-1)})}$$

with $\boldsymbol{r}^{(i-1)} = \boldsymbol{x}_{\mathrm{dirty}} - \kappa \mathrm{Re}\{\boldsymbol{\Phi}^\dagger \boldsymbol{\Phi}\}\boldsymbol{x}^{(i-1)}$

- "Series" expression for $\boldsymbol{x}^{(I)}$:

$$\boxed{\widehat{\boldsymbol{x}} \triangleq \boldsymbol{x}^{(I)} = \sum_{i=1}^{I} \mathbf{N}_{\widehat{\boldsymbol{\theta}}^{(i)}}(\boldsymbol{r}^{(i-1)}, \boldsymbol{x}^{(i-1)})}$$

- Training losses for $\mathbf{N}_{\widehat{\boldsymbol{\theta}}^{(i)}}$ sequence:

$$\boxed{\widehat{\boldsymbol{\theta}}^{(i)} = \arg\min_{\boldsymbol{\theta}^{(i)} \in \mathbb{R}^Q} \frac{1}{L} \sum_{l=1}^{L} \|\boldsymbol{x}_l^{\star} - [\boldsymbol{x}_l^{(i-1)} + \mathbf{N}_{\boldsymbol{\theta}^{(i)}}(\boldsymbol{r}_l^{(i-1)}, \boldsymbol{x}_l^{(i-1)})]_+\|_1}$$

# R2D2: a learned version of CLEAN



R2D2's high-dynamic range networks are trained from low-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:

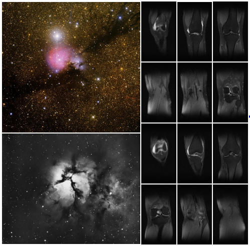

Low-dynamic range database

- ✓ artefacts removal
- ✓ size adjustment
- ✓ exponentiation

Pre-processing

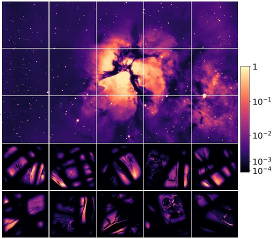


High-dynamic range database

# R2D2: a learned version of CLEAN

R2D2's high-dynamic range networks are trained from low-dynamic range databases, with flexibility in the underpinning network architecture.

- U-Net architecture underpinning the first R2D2 incarnation (learned version of Högbom CLEAN):

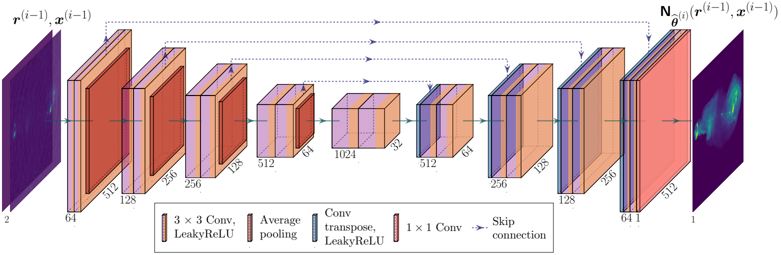

# R2D2: a learned version of CLEAN

R2D2's high-dynamic range networks are trained from low-dynamic range databases, with flexibility in the underpinning network architecture.

- R2D2-Net architecture underpinning the second R2D2 incarnation, aka R3D3 (learned version of Cotton-Schwab CLEAN):

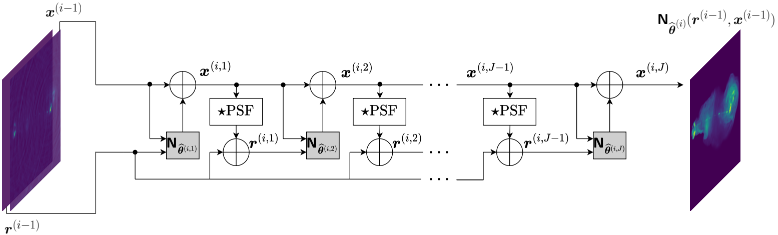

# R2D2: RI training setting



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

- Groundtruth images:
  - ✓ Exponentiated images with dynamic ranges in $[10^3, 5 \times 10^5]$
  - ✓ 10k images with $512 \times 512$ pixels
- Observation model:
  - ✓ 20k VLA sampling patterns with $[0.2, 2]$ million data points
  - ✓ Input SNR commensurate to image dynamic range
  - ✓ Briggs weighting
  - ✓ Super-resolution factor 1.5

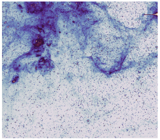

Astronomical images

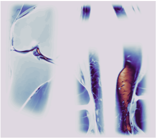

Medical images

# R2D2: RI simulation setting



Our first instance of R2D2 is validated in simulation for VLA imaging.

- Groundtruth images:
  - ✓ Real radio images with a dynamic range $[10^3, 5 \times 10^5]$
  - ✓ 4 images with $512 \times 512$ pixels
- Observation model:
  - ✓ 50 VLA sampling patterns with $[0.2, 2]$ million data points
  - ✓ Input SNR commensurate to image dynamic range
  - ✓ Briggs weighting
  - ✓ Super-resolution factor 1.5

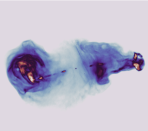

3c353

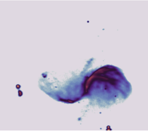

M106

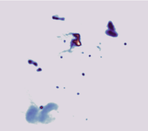

PSZ2 G165.68+44.01

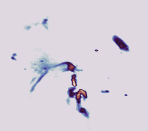

ACO 2034

# R2D2: quantitative RI simulation results

Quantitative metrics confirm R2D2 brings superior precision to AIRI and uSARA... at a fraction of the cost.

- Simulation results:

| **Model** | **SNR ± std (dB)** | **logSNR ± std (dB)** | $t_{\text{tot}} \pm \text{std (sec)}$ | iteration # |
|---|---|---|---|---|
| CLEAN | 13.6±3.6 | 10.3±3.5 | 65.9±14.2 | 9±1*₁ |
| uSARA | 30.8±1.9 | 21.9±3.3 | 4184.2±1548.9 | 1103±373 |
| AIRI | 31.3±2.3 | 21.9±4.4 | 3478.8±1531.4 | 5000±0.0*₂ |
| U-Net | 20.5±2.7 | 6.6±3.3 | 1.1±0.1 | 1 |
| R2D2-Net | 33.7±1.7 | 24.0±4.7 | 1.1±0.1 | 1 |
| R2D2 | 33.7±1.5 | 25.0±4.9 | 2.9±0.3 | 15 |
| **R3D3** | **34.0±1.6** | **25.3±4.7** | **2.2±0.3** | **8** |

$*_1$: number of "major cycles"
$*_2$: max. iteration number systematically reached

# R2D2 trained for VLA: visual simulations results



Reconstruction results from simulated observations of the 3C353 source.

- Model evolution across iterations for R2D2 incarnation with U-Net:

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

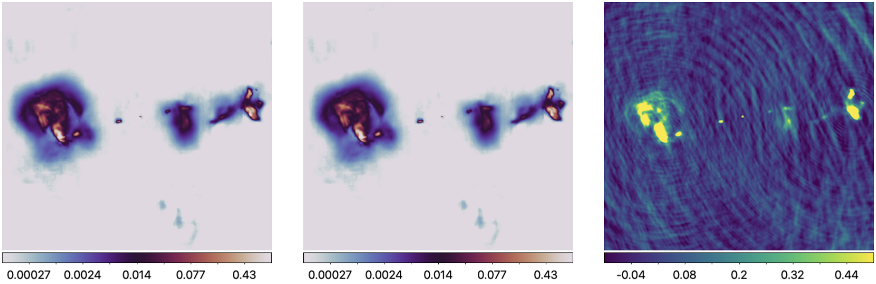


Log scale visualisation | Log scale visualisation | Linear scale visualisation

# R2D2 trained for VLA: visual simulations results



Reconstruction results from simulated observations of the 3C353 source.

- Model evolution across iterations for R2D2 incarnation with U-Net:

$\boldsymbol{x}^{(2)}$ $\quad$ $\mathbf{N}_{\widehat{\boldsymbol{\theta}}^{(2)}}(\boldsymbol{r}^{(1)}, \boldsymbol{x}^{(1)})$ $\quad$ $\boldsymbol{r}^{(1)}$

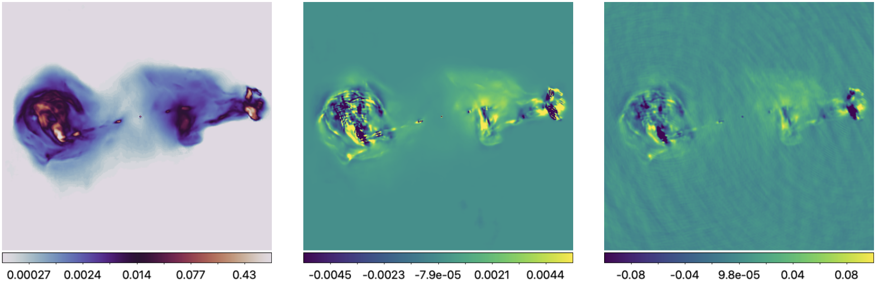


Log scale visualisation $\quad$ Linear scale visualisation $\quad$ Linear scale visualisation

# R2D2 trained for VLA: visual simulations results

Reconstruction results from simulated observations of the 3C353 source.

- Model evolution across iterations for R2D2 incarnation with U-Net:

$\boldsymbol{x}^{(15)}$

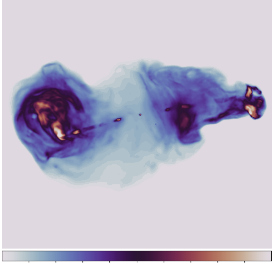


Log scale visualisation

$\mathbf{N}_{\widehat{\boldsymbol{\theta}}^{(15)}}(\boldsymbol{r}^{(14)}, \boldsymbol{x}^{(14)})$

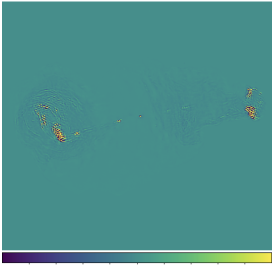


Linear scale visualisation

$\boldsymbol{r}^{(14)}$

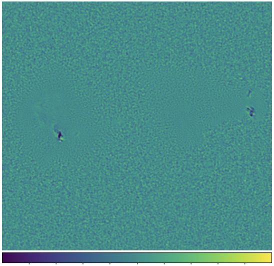


Linear scale visualisation

# R2D2: real imaging quiz



Which methods hide behind these Cygnus A images formed from real VLA observations?

- **Dirty image** (2.05GHz, 20MB data, 512 x 512 pixels):

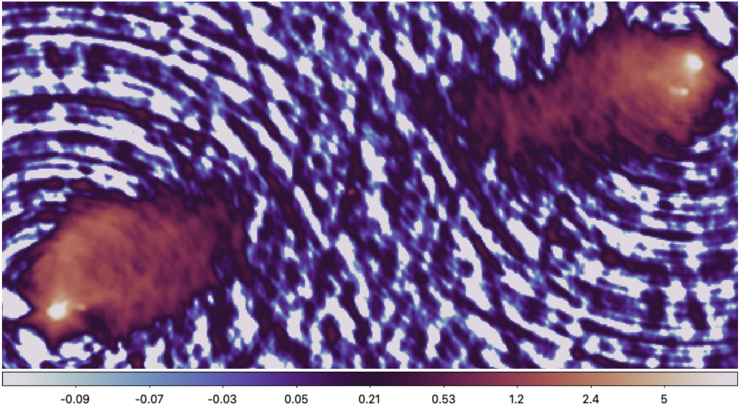

# R2D2: real imaging quiz



Which methods hide behind these Cygnus A images formed from real VLA observations?

- **Method #1** (2.05GHz, 20MB data, 512 x 512 pixels):

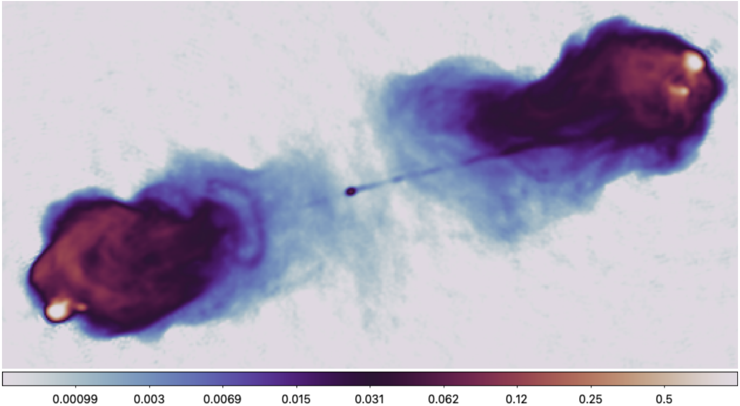

# R2D2: real imaging quiz



Which methods hide behind these Cygnus A images formed from real VLA observations?

- **Method #2** (2.05GHz, 20MB data, 512 x 512 pixels):

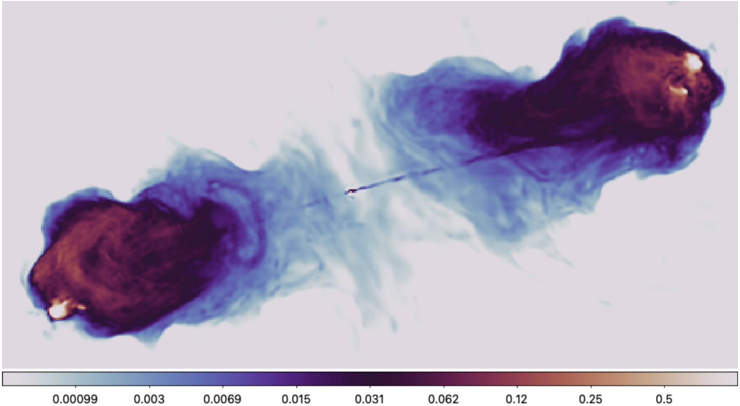

# R2D2: real imaging quiz

Which methods hide behind these Cygnus A images formed from real VLA observations?

- **Method #1: CLEAN**

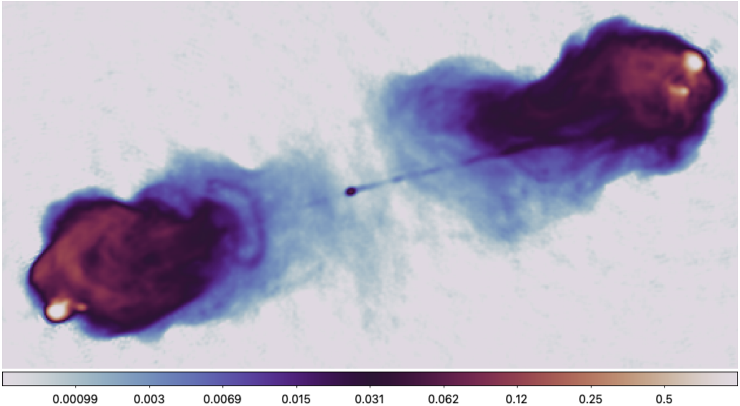

# R2D2: real imaging quiz



Which methods hide behind these Cygnus A images formed from real VLA observations?

- **Method #2: R3D3**

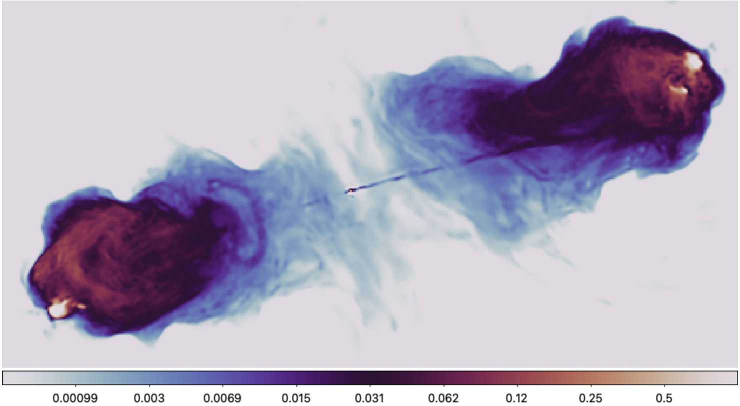

# On the R2D2 model uncertainty

Generating "R2D2 samples" from multiple series, trained from different random initialisations, enables tracking model uncertainty.

- In few iterations only, model uncertainty decreases to very low levels:

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(i)})$

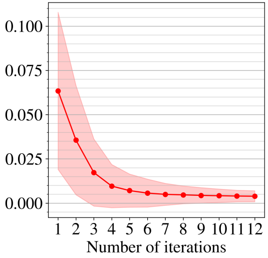


$\mu(\bar{\boldsymbol{x}}^{(1)})$

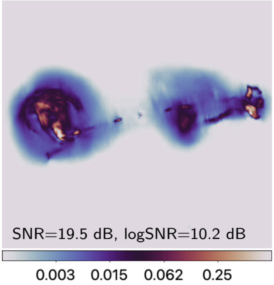


Log scale visualisation

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(1)})$

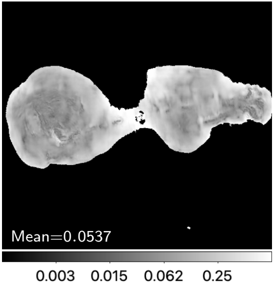


Log scale visualisation

# On the R2D2 model uncertainty



Generating "R2D2 samples" from multiple series, trained from different random initialisations, enables tracking model uncertainty.

- In few iterations only, model uncertainty decreases to very low levels:

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(i)})$

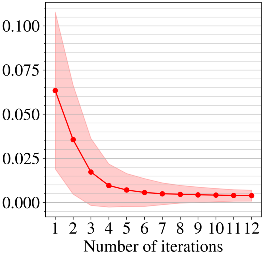


$\mu(\bar{\boldsymbol{x}}^{(3)})$

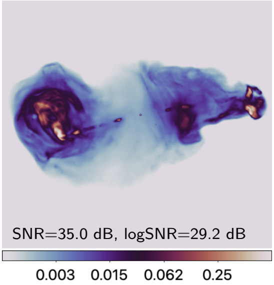


Log scale visualisation

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(3)})$

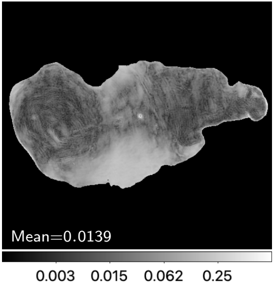


Log scale visualisation

# On the R2D2 model uncertainty



Generating "R2D2 samples" from multiple series, trained from different random initialisations, enables tracking model uncertainty.

- In few iterations only, model uncertainty decreases to very low levels:

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(i)})$

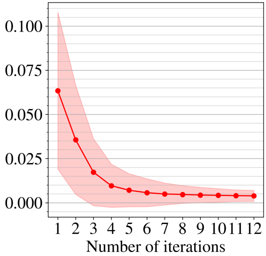


$\mu(\bar{\boldsymbol{x}}^{(12)})$

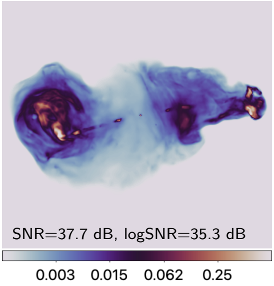


Log scale visualisation

$[\sigma/\mu](\bar{\boldsymbol{x}}^{(12)})$

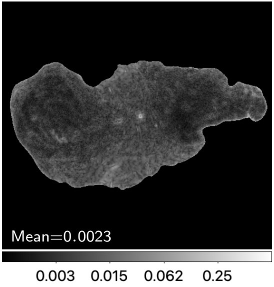


Log scale visualisation

# On the transfer of technology to MRI



R2D2 transfers seamlessly to MRI, enabling accurate reconstruction from accelerated radial k-space sampling with data-intensive multi-coil acquisitions.

- R2D2 supersedes scalable competitors up to high acceleration:

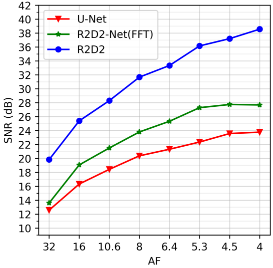


SNR vs Acceleration Factor (AF)

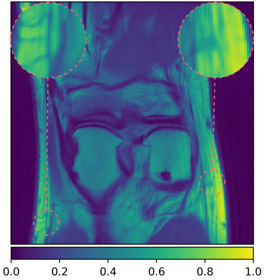


Ground Truth magnitude

SNR=16.60 dB

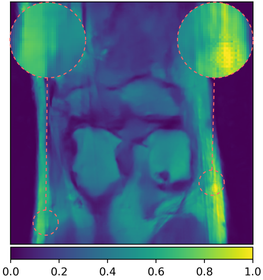


U-Net (AF=16, 29 coils)

# On the transfer of technology to MRI

R2D2 transfers seamlessly to MRI, enabling accurate reconstruction from accelerated radial k-space sampling with data-intensive multi-coil acquisitions.

- R2D2 supersedes scalable competitors up to high acceleration:

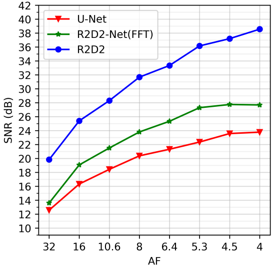


SNR vs Acceleration Factor (AF)

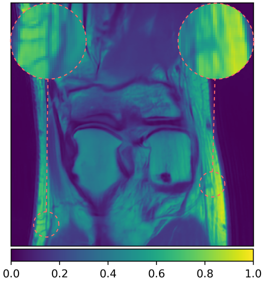


Ground Truth magnitude

SNR=17.54 dB

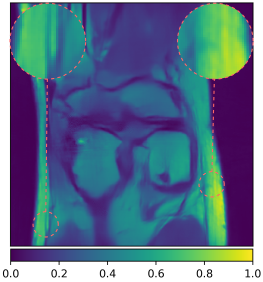


R2D2-Net(FFT) (AF=16, 29 coils)

# On the transfer of technology to MRI



R2D2 transfers seamlessly to MRI, enabling accurate reconstruction from accelerated radial k-space sampling with data-intensive multi-coil acquisitions.

- R2D2 supersedes scalable competitors up to high acceleration:

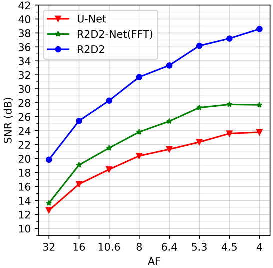


SNR vs Acceleration Factor (AF)

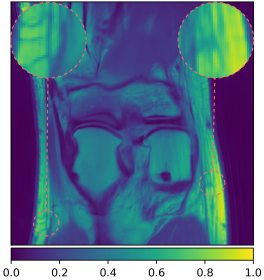


Ground Truth magnitude

SNR=22.25 dB

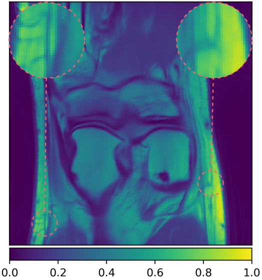


R2D2 (AF=16, 29 coils)

# Conclusion & future work: more seriously



R2D2 offers a new regime of quality and speed in large-scale high-resolution high-dynamic range computational imaging in radio astronomy (but not only), paving the way towards ultra-fast acquisition and reconstruction.

**Upcoming evolutions:**

- ☛ Investigate R2D2 convergence and generalisability
- ☛ Add calibration and polarisation functionality

**Python and Matlab code**

- ☛ Enjoy BASP's new Computational Imaging Library: BASPLib

Anything on black holes?

# Oops... yes... something on black holes



A bespoke R2D2-Net might come to light...

Ground truth

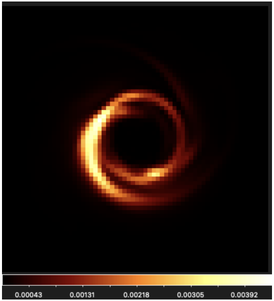


Linear scale visualisation

uSARA

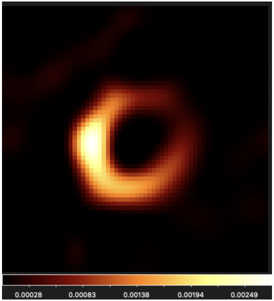


Linear scale visualisation

R2D2-Net

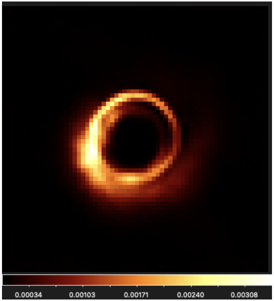


Linear scale visualisation