How to clean deep and fast with R2D2

- unveiling a Star Wars hidden fact -

Prof. Yves Wiaux, Edinburgh on behalf of BASP & Co.

Swiss SKA Days 2023, Zurich Switzerland





• SYNTHESIS IMAGING IN RADIO ASTRONOMY From CLEAN to uSARA & AIRI algorithms

• R2D2 ALGORITHM: A newborn DNN series approach





SYNTHESIS IMAGING IN RADIO ASTRONOMY From CLEAN to uSARA & AIRI algorithms

- Repetti & Wiaux, SIOPT, 2019, arXiv:1907.11486
 - Pesquet et al., SIIMS, 2020, arXiv:2012.13247
 - Terris et al., MNRAS, 2022, arXiv:2202.12959
 - Dabbech et al., ApJL, 2022, arXiv:2207.11336
- Wilber et al., MNRAS, 2023a,b, arXiv:2302.14148,9





Aperture synthesis by radio interferometry

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

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

$$y = \Phi \bar{x} + n$$

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

$$oldsymbol{y}
ightarrow oldsymbol{ar{x}}$$

Accurate image models are needed for precision (*i.e.* high resolution and high dynamic range)



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 \bar{x} with PSF Re{ $\Phi^{\dagger}\Phi$ } δ :

$$\mathbf{x}_{ ext{dirty}} = \kappa \operatorname{\mathsf{Re}}\{\mathbf{\Phi}^{\dagger}\mathbf{y}\} \simeq \kappa \operatorname{\mathsf{Re}}\{\mathbf{\Phi}^{\dagger}\mathbf{\Phi}\}\boldsymbol{\delta} \star \bar{\mathbf{x}} + \mathbf{n}'$$

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

CLEAN iteration structure (peeling bright sources in residual image):

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

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

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





Powerful optimisation framework

Convex optimisation provides a powerful framework to solve inverse problems via iterative algorithms.

Image reconstruction formulated as a convex optimisation problem:

$$\mathbf{x}^{\star} \in \operatorname*{argmin}_{\mathbf{x}} \Big\{ g(\mathbf{x}; \mathbf{y}) = f(\mathbf{x}; \mathbf{y}) + r(\mathbf{x}) \Big\}$$

with $f(\mathbf{x}; \mathbf{y})$: data-fidelity term; $r(\mathbf{x})$: regularisation term

- The theory is robust and versatile:
 - $\checkmark\,$ algorithms endowed with convergence guarantees
 - $\checkmark\,$ advanced regularisation enables precision
 - ✓ parallel algorithmic structures enable scalability





Forward-Backward algorithm

The Forward-Backward (FB) algorithm is a simple and flexible optimisation structure.

FB minimisation task:

$$\mathbf{x}^{\star} \in \operatorname*{argmin}_{\mathbf{x}} \Big\{ g(\mathbf{x}; \mathbf{y}) = f(\mathbf{x}; \mathbf{y}) + r(\mathbf{x}) \Big\}$$

with f(x; y): differentiable; r(x): differentiable or not

► FB iteration structure:

$$\boldsymbol{x}^{(i)} = \operatorname{prox}_{\gamma r} \left(\boldsymbol{x}^{(i-1)} - \gamma \nabla f \left(\boldsymbol{x}^{(i-1)} \right) \right)$$

with $\gamma < 2/L(\nabla f)$

- ✓ forward gradient-descent data-fidelity step
- \checkmark backward regularisation step involving $\mathrm{prox}_{\gamma r}$
- \checkmark the proximal operator $\mathrm{prox}_{\gamma r}$ is an image denoiser





Unconstrained SARA (uSARA)

Unconstrained SARA leverages FB with handcrafted regularisation for monochromatic intensity imaging.

• Data fidelity term: $f(\mathbf{x}, \mathbf{y}) = ||\mathbf{y} - \mathbf{\Phi}\mathbf{x}||_2^2$ (Gaussian noise)

Regularisation term: log-sum prior (generalising l₁) promoting average sparsity in a redundant wavelet dictionary

$$\mathbf{r}(\mathbf{x}) = \eta \sum_{b=1}^{B} \rho \log \left(1 + \rho^{-1} \left| \left(\mathbf{\Psi}^{\dagger} \mathbf{x} \right)_{b} \right| \right) + \iota_{\mathbb{R}^{N}_{+}}(\mathbf{x})$$

uSARA iteration structure:

$$\boldsymbol{x}^{(i)} = \operatorname{prox}_{\gamma r} \left(\boldsymbol{x}^{(i-1)} + \gamma \operatorname{Re} \{ \boldsymbol{\Phi}^{\dagger} \boldsymbol{y} \} - \gamma \operatorname{Re} \{ \boldsymbol{\Phi}^{\dagger} \boldsymbol{\Phi} \} \boldsymbol{x}^{(i-1)} \right)$$

with $\gamma < 2/\|\text{Re}\{\pmb{\Phi}^{\dagger}\pmb{\Phi}\}\|_{\mathrm{S}}$

uSARA's proximal operator is iterative, affecting scalability





AI for Regularisation in Imaging (AIRI)

AIRI leverages FB, plugging a learned DNN denoiser in lieu of a proximal operator for monochromatic intensity imaging, in a plug-and-play (PnP) approach.

- ► Data fidelity term: $f(\mathbf{x}, \mathbf{y}) = ||\mathbf{y} \mathbf{\Phi}\mathbf{x}||_2^2$ (Gaussian noise)
- Regularisation term: implicitly defined by a learned DNN denoiser J
- AIRI iteration structure:

$$\mathbf{x}^{(i)} = \mathbf{J} \left(\mathbf{x}^{(i-1)} + \gamma \operatorname{Re} \{ \mathbf{\Phi}^{\dagger} \mathbf{y} \} - \gamma \operatorname{Re} \{ \mathbf{\Phi}^{\dagger} \mathbf{\Phi} \} \mathbf{x}^{(i-1)} \right)$$

with $\gamma < 2/\|\text{Re}\{\pmb{\Phi}^{\dagger}\pmb{\Phi}\}\|_{\text{S}}$

- Learning opens the door to powerful physical regularisation
- DNNs provides acceleration over sub-iterative proximal operators
- ▶ Denoisers blind to **Φ**, but requires a tailored training approach...





AIRI & uSARA for ASKAP imaging

"Dancing ghosts" reconstruction from wide-field ASKAP data.

CLEAN (0.9GHz, 7GB data, 4k × 4k pixels, 58 CoreH):









AIRI & uSARA for ASKAP imaging

"Dancing ghosts" reconstruction from wide-field ASKAP data.

▶ uSARA (0.9GHz, 7GB data, 4k × 4k pixels, 770 CoreH):









AIRI & uSARA for ASKAP imaging

"Dancing ghosts" reconstruction from wide-field ASKAP data.

► AIRI (0.9GHz, 7GB data, 4k × 4k pixels, 203 CoreH):









AIRI & uSARA algorithms are highly iterative, which fundamentally limits their scalability.





A scalability challenge remains

End-to-end DNNs, though suffering from generalisation issues, remain appealing as they provide almost real-time reconstruction from the dirty image.



DNN input:

$$\mathbf{x}_{ ext{dirty}} = \kappa \operatorname{\mathsf{Re}}\{\mathbf{\Phi}^{\dagger}\mathbf{y}\}$$

with $\kappa = 1/\max(\mathsf{Re}\{\mathbf{\Phi}^{\dagger}\mathbf{\Phi}\}m{\delta})$

DNN output: x^*

$$\textbf{\textit{x}}^{\star} = \textbf{\textit{G}}(\textbf{\textit{x}}_{\mathrm{dirty}})$$

Existing DNNs do not match the precision of uSARA or AIRI.





R2D2 ALGORITHM: from DNN series to astronomical imaging

Aghabiglou et al., Proc. ICASSP 2023, arXiv:2210.1606

Aghabiglou et al., ApJL, in prep.

Aghabiglou et al., MNRAS, in prep.





Star Wars hidden fact

Did you know R2D2 stands for Residual-to-Residual DNN series for high-Dynamic range imaging?



Star Wars: Episode I – The Phantom Menace (1999)





R2D2: a learned version of matching pursuit

R2D2 applies end-to-end DNNs iteratively, each network taking the residual dirty image of the previous iteration as an input, and reconstructing the residual between the ground truth and the reconstruction of the previous iteration.

► R2D2 iteration structure:

$$\boldsymbol{x}^{(i)} = \boldsymbol{x}^{(i-1)} + \boldsymbol{\mathsf{G}}^{(i)}\left(\boldsymbol{r}^{(i-1)}
ight)$$

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

• "Series" expression for
$$x^{(l)}$$
:

$$\mathbf{x}^{\star} = \mathbf{x}^{(l)} = \sum_{i=1}^{l} \mathbf{G}^{(i)}(\mathbf{r}^{(i-1)})$$

► Training losses for **G**^(*i*) sequence:

$$\min_{\theta} \frac{1}{L} \sum_{\ell=1}^{L} \| \mathbf{G}_{\theta}^{(i)}(\mathbf{r}_{\ell}^{(i-1)}) + \mathbf{x}_{\ell}^{(i-1)} - \bar{\mathbf{x}}_{\ell} \|_{1}$$





R2D2: a learned version of matching pursuit

R2D2's high-dynamic range networks are trained from low-dynamic range databases, with an advanced network architecture.

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







R2D2: a learned version of matching pursuit

R2D2's high-dynamic range networks are trained from low-dynamic range databases, with an advanced network architecture.

New U-WDSR architecture combining U-Net and WDSR:







R2D2 trained for VLA: quantitative simulation results 16/19

Quantitative metrics confirm R2D2 brings the same precision as AIRI and uS-ARA... at a fraction of the cost.

Model	$\mathbf{SNR}(dB)$	$\log SNR(dB)$	time(s.)	iteration $\#$
CLEAN	$13.3 {\pm} 0.1$	14.3±0.4	99±4	$4.8 \pm 0.2^{*_1}$
uSARA	$25.3{\pm}0.2$	$25.5 {\pm} 0.1$	$4419{\pm}111$	2299±34
AIRI	27.4 ± 0.3	28.7 ± 0.2	$3349{\pm}159$	6000* ²
U-WDSR	$18.5 {\pm} 0.3$	11.5±0.4	$1.0{\pm}0.2$	1
R2D2	26.8±0.4	27.3±0.3	$9.8 {\pm} 0.6$	10 *3

Simulation results:

Table: Average metrics and 95% confidence intervals, over a test dataset of 150 inverse problems. Training and test are for 512×512 images, VLA sampling with [0.2,2] million data points, Briggs weighting, with varying noise level and super-resolution factor 1.5.

*1: Number of "major cycles"

*2: Maximum iteration number systematically reached

 $*_3$: After $x^{(4)}$ the results are generated by repeating $G^{(4)}$ multiple times.





R2D2 trained for VLA: visual simulations results

Reconstruction results from simulated observations of the 3C353 source.

Model evolution across iterations:







R2D2 trained for VLA: visual simulations results

Reconstruction results from simulated observations of the 3C353 source.

Model evolution across iterations:



Log scale visualisation

Linear 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:



Linear scale visualisation

Linear scale visualisation



Cygnus A reconstruction results from real VLA observations.

▶ Dirty image (2.05GHz, 20MB data, 512 × 512 pixels):







Cygnus A reconstruction results from real VLA observations.

CLEAN (2.05GHz, 20MB data, 512 × 512 pixels):







Cygnus A reconstruction results from real VLA observations.

▶ **uSARA** (2.05GHz, 20MB data, 512 × 512 pixels):







Cygnus A reconstruction results from real VLA observations.

► AIRI (2.05GHz, 20MB data, 512 × 512 pixels):







Cygnus A reconstruction results from real VLA observations.

► U-WDSR (2.05GHz, 20MB data, 512 × 512 pixels):



0.00052 0.0015 0.0035 0.0074 0.015 0.031 0.062 0.13 0.25





Cygnus A reconstruction results from real VLA observations.

▶ **R2D2** (2.05GHz, 20MB data, 512 × 512 pixels):







Hybrid algorithms at the interface of optimisation and deep learning can offer a new regime of quality and speed in large-scale high-resolution high-dynamic range computational imaging in radio astronomy, possibly paving the way towards near-real time imaging.

Ongoing and future AIRI and R2D2 evolutions:

- Investigate advanced AIRI and R2D2 DNN architectures and losses
- Investigate R2D2 convergence, robustness, generalisability
- Add calibration, uncertainty quantification, and other functionalities
- Translate current Matlab code into C++ (Puri-Psi on GitHub)



