Learned Image Reconstruction for Interferometric Imaging

Matthijs Mars*, Marta Betcke, Jason McEwen

* matthijs.mars.20@ucl.ac.uk

Mullard Space Science Laboratory (MSSL), University College London (UCL)

September 28, 2023

Outline

1. Interferometric Imaging
2. The Inverse Problem
3. Learned Image Reconstruction
4. Reconstruction for SPIDER
5. Reconstruction for Radio telescopes
6. Conclusion

Interferometric Imaging

Interferometric Measurements

The baselines measure the visibilities $y(u, v)$ which correspond to a Fourier measurement of the image:

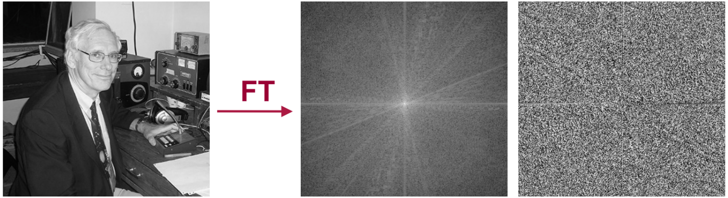

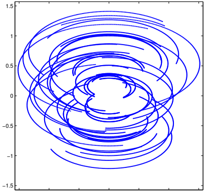

Credit: NRAO

SPIDER

Segmented Planar Imaging Detector for Electro-optical Reconnaissance

- SPIDER is new interferometric optical imaging device developed by UC Davis and Lockheed Martin.
- Lenslet array to measure multiple interferometric baselines and photonic integrated circuits (PICs) for miniaturization.
- Reduces weight, cost and power consumption of optical telescopes.

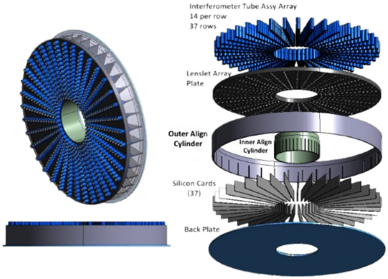

Kendrick et al. 2013

SPIDER

Challenges for SPIDER:

- ▷ Sparse sampling of Fourier measurements. ($M = 4440$)
- ▷ Aliasing because of sampling distribution.
- ▷ Real-time reconstruction.

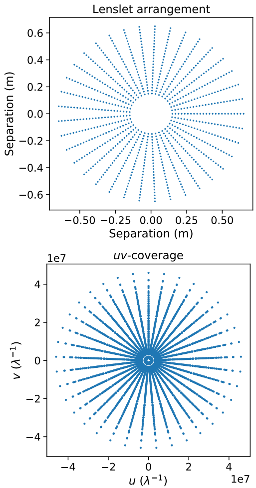

Radio Telescopes

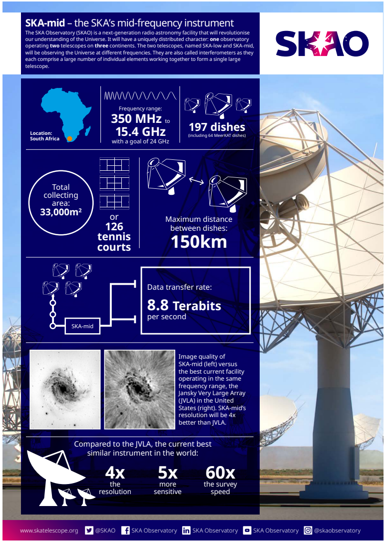

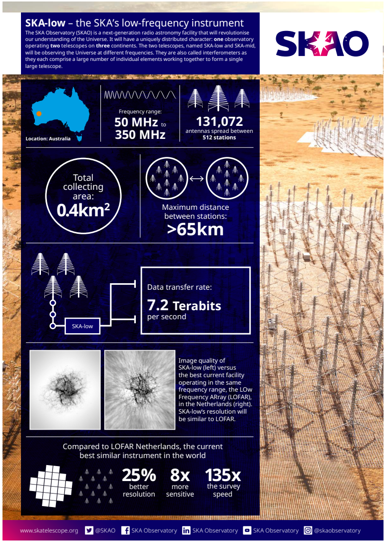

Radio Telescopes

Challenges for Radio telescopes:

- ▷ Large amount of visibilities. ($M \gg N$)
- ▷ Visibility coverage depends on pointing, earth rotation, etc.

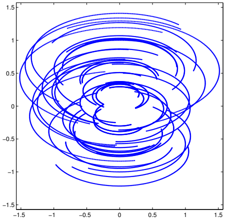

The Inverse Problem

Inverse Problem

We can write this particular problem as an inverse problem

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

with

- $\boxed{\boldsymbol{y} \in \mathbb{C}^M}$ – Non-uniformly distributed Fourier measurements at (u_j, v_j) for $j = 0, \dots, M$
- $\boxed{\boldsymbol{\Phi} : \mathbb{R}^N \to \mathbb{C}^M}$ – The forward operation of our measurement operator of the telescope
- $\boxed{\boldsymbol{x} \in \mathbb{R}^N}$ – The image to observe with N pixels
- $\boxed{\boldsymbol{n} \in \mathbb{C}^M}$ – Additive complex Gaussian noise

Variational Regularisation

Variational regularisation: Find $\boldsymbol{x}^\star$ that satisfies

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x} \in \mathbb{R}^N}{\arg\min}\; \mathcal{L}(\boldsymbol{\Phi} \boldsymbol{x}, \boldsymbol{y}) + \lambda \mathcal{R}(\boldsymbol{x}),$$

which we would typically solve in astronomical interferometry problems as

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x} \in \mathbb{R}^N}{\arg\min}\; ||\boldsymbol{\Phi} \boldsymbol{x} - \boldsymbol{y}||_2^2 + \lambda ||\boldsymbol{\Psi}^\dagger \boldsymbol{x}||_1,$$

where $\boldsymbol{\Phi}$ models the measurement operator and $\boldsymbol{\Psi}$ is typically a dictionary of wavelet bases.

Learned Image Reconstruction

Learned Image Reconstruction

Using deep learning to solve the inverse problem:

- Data-driven priors
- Higher reconstruction quality
- Better computational efficiency

Learned Reconstruction Approaches

Learned solvers are approaches where the parameters θ are Learned from training data.

Learned regularisation: Learn a regulariser $\boxed{\mathcal{R}_\theta}$ and solve

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x} \in \mathbb{R}^N}{\arg\min}\ \mathcal{L}(\Phi \boldsymbol{x}, \boldsymbol{y}) + \boxed{\lambda \mathcal{R}_\theta(\boldsymbol{x})}$$

Learned Reconstruction Approaches

Learned solvers are approaches where the parameters θ are Learned from training data.

Learned regularisation: Learn a regulariser $\boxed{\mathcal{R}_\theta}$ and solve

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x}\in\mathbb{R}^N}{\arg\min}\ \mathcal{L}(\boldsymbol{\Phi}\boldsymbol{x}, \boldsymbol{y}) + \boxed{\lambda\mathcal{R}_\theta(\boldsymbol{x})}$$

Learned Sequential models: Learn a sequential Pseudo-inverse $\boxed{\boldsymbol{\Phi}_\theta^\dagger}$ such that

$$\boldsymbol{x}^\star = \boldsymbol{\Phi}_\theta^\dagger\boldsymbol{y}, \quad \boxed{\boldsymbol{\Phi}_\theta^\dagger = \boldsymbol{B}_\theta \circ \boldsymbol{\Phi}^\dagger \circ \boldsymbol{C}_\theta}$$

Learned Reconstruction Approaches

Learned solvers are approaches where the parameters θ are Learned from training data.

Learned regularisation: Learn a regulariser $\boxed{\mathcal{R}_\theta}$ and solve

$$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x}\in\mathbb{R}^N} \mathcal{L}(\Phi\boldsymbol{x}, \boldsymbol{y}) + \boxed{\lambda\mathcal{R}_\theta(\boldsymbol{x})}$$

Learned Sequential models: Learn a sequential Pseudo-inverse $\boxed{\Phi_\theta^\dagger}$ such that

$$\boldsymbol{x}^\star = \Phi_\theta^\dagger\boldsymbol{y}, \quad \boxed{\Phi_\theta^\dagger = \boldsymbol{B}_\theta \circ \Phi^\dagger \circ \boldsymbol{C}_\theta}$$

Learned Unrolled Iterative Algorithms: Learn how to iteratively update the reconstruction

$$\boxed{\boldsymbol{x}_{\boldsymbol{i}+1} = \Lambda_{\boldsymbol{\theta}}(\boldsymbol{x}_{\boldsymbol{i}}, \nabla\mathcal{L}(\Phi\boldsymbol{x}_{\boldsymbol{i}}, \boldsymbol{y})), \quad \text{for i in } (0, \ldots, \text{N})}$$

Learned Post-processing

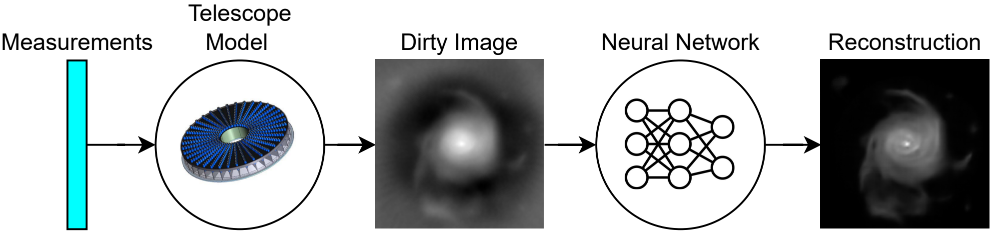

Learned Post-processing Network (U-Net)

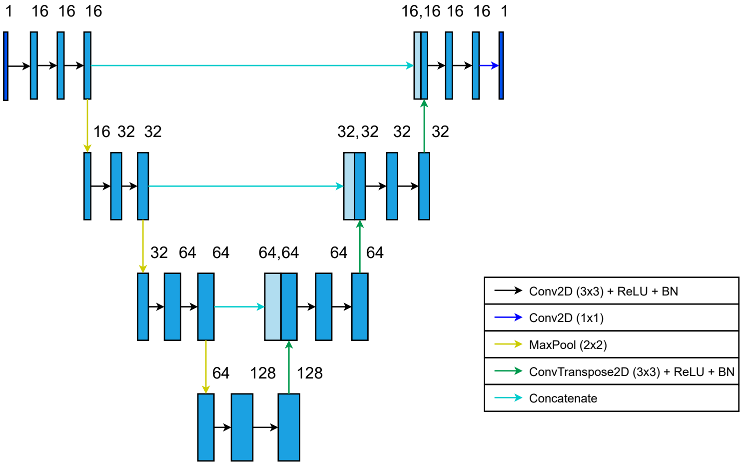

- ▷ Input: Dirty image
- ▷ Multi-resolution denoising

Learned Unrolled Iterative approach

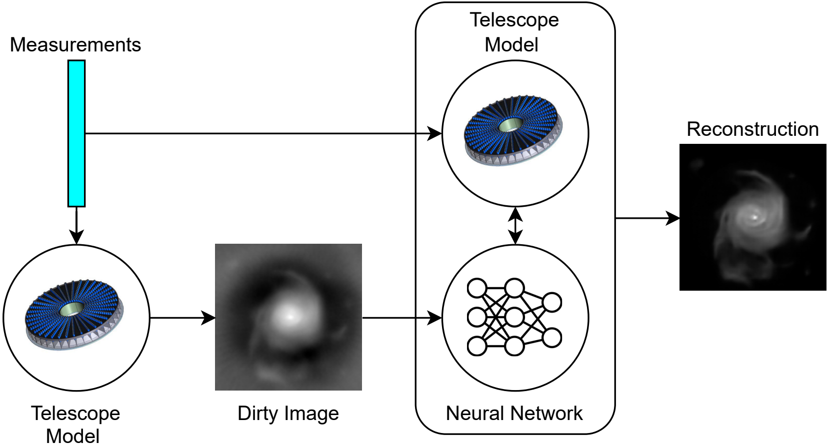

GU-Net Architecture

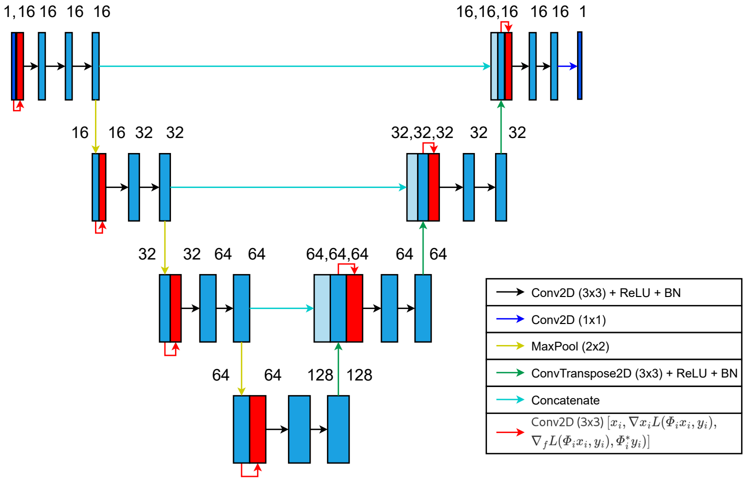

Adding in the gradient of the data fidelity term to the U-Net structure.

For ℓ_2-loss:

$$\mathcal{L} = ||\boldsymbol{\Phi} \boldsymbol{x} - \boldsymbol{y}||_{\ell_2}^2,$$

we calculate the gradient as follows:

$$\nabla \mathcal{L}(\boldsymbol{\Phi} \boldsymbol{x}, \boldsymbol{y}) \propto \boldsymbol{\Phi}^\dagger (\boldsymbol{\Phi} \boldsymbol{x} - \boldsymbol{y})$$

Learned Unrolled Iterative approach

On each scale of the U-Net structure, we apply a subsampling of the measurement operator to use for the gradient of the data fidelity term.

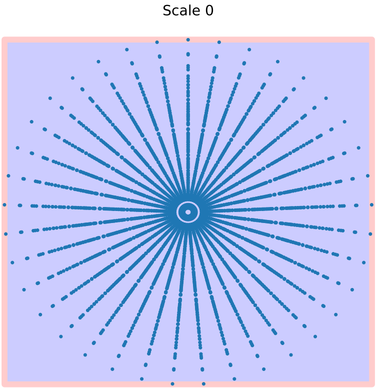

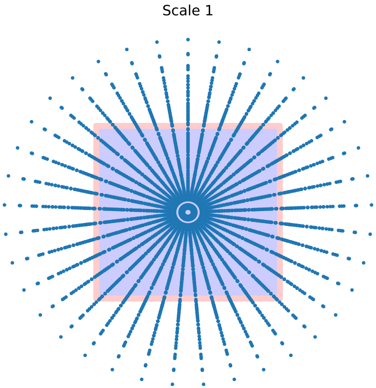

Reconstruction for SPIDER

Reconstructed SPIDER images

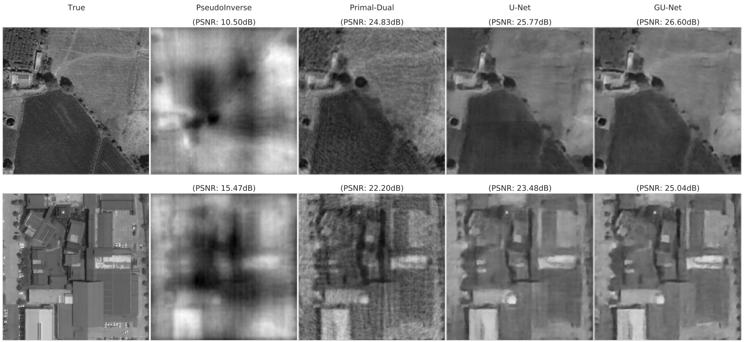

- ▷ Imaging time speed-up of 50-600× relative to classical approaches.
- ▷ Dramatic reduction in computational time opens up real time imaging with SPIDER for the first time.

Computational Efficiency

Name	Operator evaluations	Average reconstruction time (ms)	Training time (mins)
Pseudo-inverse (1 GPU)	1	5.50	-
U-Net (1 GPU)	1	10.7	$\sim$30
GU-Net (1 GPU)	11*	42.1	$\sim$100
Primal-Dual (300its, 1 CPU)	600	4.7×10^4	-

*Refers to operator evaluation at the finest scale, which dominates the computational time of the GU-Net.

- Imaging time speed-up of 50-600$\times$ relative to classical approaches.
- Dramatic reduction in computational time opens up real time imaging with SPIDER for the first time.

Reconstruction for Radio telescopes

Challenge in Radio Interferometry

Challenge: The visibility coverage is different for every observation

We compare:

- ▷ **True coverage (oracle)**:
  Actual visibility coverage, full retraining required for every observation.

Challenge in Radio Interferometry

Challenge: The visibility coverage is different for every observation

We compare:

- ▷ **True coverage (oracle)**:
  Actual visibility coverage, full retraining required for every observation.
- ▷ **Single coverage**:
  Different visibility coverage, no retraining required for every observation.

Challenge in Radio Interferometry

Challenge: The visibility coverage is different for every observation

We compare:

- **True coverage (oracle)**:
  Actual visibility coverage, full retraining required for every observation.
- **Single coverage**:
  Different visibility coverage, no retraining required for every observation.
- **Distribution of coverages**:
  Distribution of visibility coverages, no retraining required for every observation.

Challenge in Radio Interferometry

Challenge: The visibility coverage is different for every observation

We compare:

- **True coverage (oracle)**:
  Actual visibility coverage, full retraining required for every observation.
- **Single coverage**:
  Different visibility coverage, no retraining required for every observation.
- **Distribution of coverages**:
  Distribution of visibility coverages, no retraining required for every observation.
- **Transfer Learning**:
  Fine-tuning through transfer learning using the **true coverage**, small amount of retraining.
  - **Single coverage transfer**
  - **Distribution of coverages transfer**

Challenge in Radio Interferometry

Challenge: The visibility coverage is different for every observation

We compare:

- ▷ **True coverage**
- ▷ **Single coverage**
- ▷ **Single coverage transfer**
- ▷ **Distribution of coverages**
- ▷ **Distribution of coverages transfer**

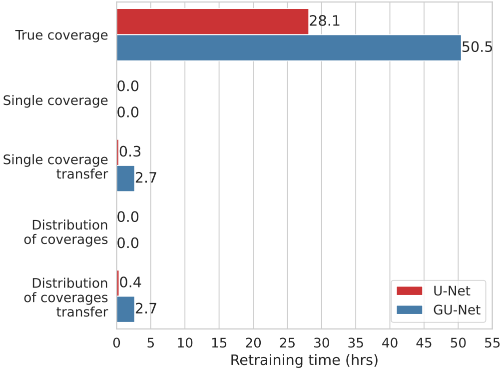

Distribution of radio interferometric reconstruction quality

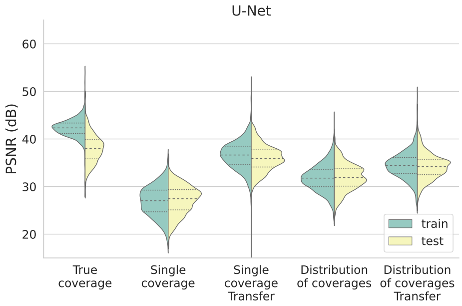

Reconstruction quality (PSNR ↑) for different training strategies.

Distribution of radio interferometric reconstruction quality

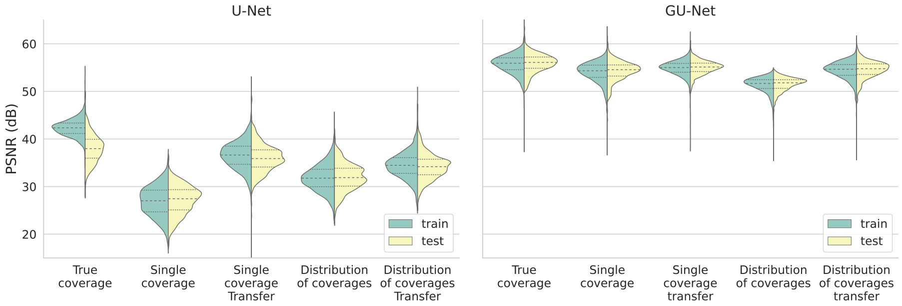

Reconstruction quality (PSNR $\uparrow$) for different training strategies.

- Superior reconstruction quality by integrating physical model of instrument and more robust to measurement operator variability.
- Imaging time speed-up of 50-600$\times$ relative to classical approaches.

Reconstructed radio interferometric images

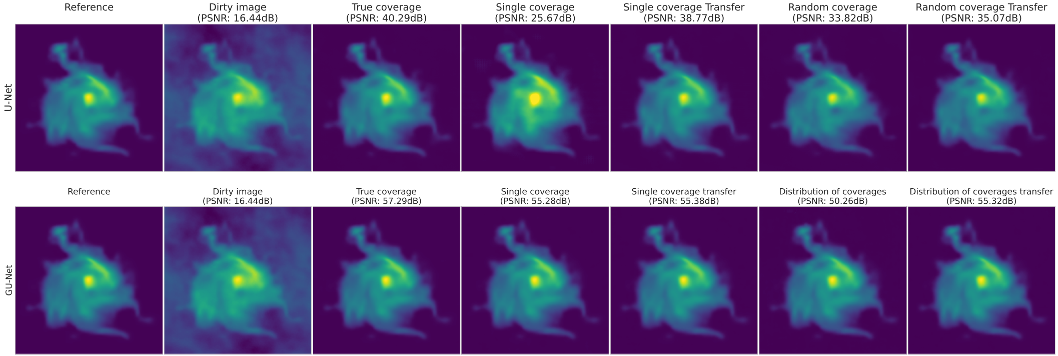

Reconstructed radio interferometric images

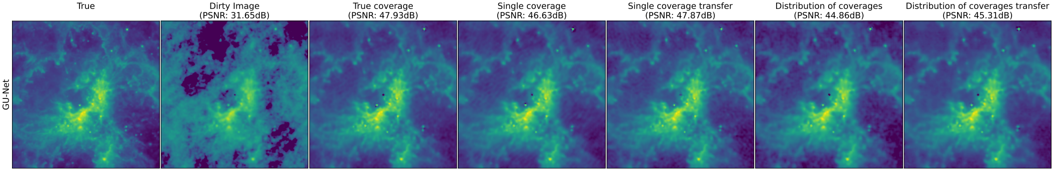

- Full end-to-end learning for radio interferometric imaging with support for varying measurement operators for the first time.

Conclusion

Conclusion

- Imaging time speed-up of 50-600$\times$ relative to classical approaches.
- Superior reconstruction quality by integrating physical model of instrument
- Full end-to-end learning for radio interferometric imaging with support for varying measurement operators for the first time.

Papers:

Mars et al. 2023, "Learned Interferometric Imaging for the SPIDER Instrument", arXiv:2301.10260

Mars et al. (in prep.), "Learned radio interferometric imaging for varying visibility coverages"