Learned Radio Interferometric Imaging for Varying Visibility Coverage

Matthijs Mars, Marta M. Betcke, Jason D. McEwen

Mullard Space Science Laboratory, University College London (UCL)

1. Overview

Radio interferometric telescopes measure visibilities (Fourier coefficients) of the sky. The distribution of the acquired Fourier samples depends on the configuration of baselines, the pointing, and the integration time of the radio telescope. This results in a visibility coverage that is unique for every observation.

Traditional imaging approaches use handcrafted priors with iterative optimization algorithms to reconstruct an image from the Fourier measurements. Using **data-driven priors**, images can be reconstructed with higher reconstruction quality while reducing the computational time (e.g. [1]).

Existing learned approaches consider a fixed measurement operator that does not change for every observation. This work addresses the challenge of training networks to support varying visibility coverages in a manner so that they do not need to be completely retrained to match the visibility coverage of every single observation.

2. Interferometric Imaging Problem

The interferometric imaging problem can be concisely described as

$$\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x} + \boldsymbol{n}, \tag{1}$$

- Non-uniformly distributed Fourier measurements, $\boldsymbol{y} \in \mathbb{C}^K$
- Measurement operator, $\boldsymbol{\Phi} : \mathbb{R}^N \to \mathbb{C}^K$
- Image, $\boldsymbol{x} \in \mathbb{R}^N$
- Measurement noise, $\boldsymbol{n} \in \mathbb{C}^K$

The measurement operator is modelled using a **non-uniform fast Fourier Transform (NUFFT, [2])**. The implementation of the NUFFT is described in [1].

Modern classical reconstruction approaches use iterative solvers to find

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x}\in X}{\arg\min} \quad \|\boldsymbol{\Phi}\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2}^2 + \lambda \|\boldsymbol{\Psi}^\dagger \boldsymbol{x}\|_{\ell_1}, \tag{2}$$

- $\mathcal{L} = \|\boldsymbol{\Phi}\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2}^2$, data fidelity term
- $\mathcal{S} = \|\boldsymbol{\Psi}^\dagger \boldsymbol{x}\|_{\ell_1}$, sparsity prior, with $\boldsymbol{\Psi}$ typically a dictionary of wavelet bases

These approaches are
- Computationally expensive, as they evaluate the measurement operator every iteration
- Limited by the prior information captured in the handcrafted prior ($\mathcal{S}$)

3. Learned Interferometric Imaging for varying Visibility Coverage

Learned imaging methods can be used to reduce computational cost and get increased reconstruction quality through data-driven priors. We use two network architecture that we proposed in [1]:
- The **U-Net**: a learned post-processing method that uses the U-Net architecture [3] to denoise the dirty image
- The **Gradient U-Net (GU-Net)**: a learned unrolled iterative method that combines a U-Net architecture with updates based on the gradient of the likelihood function, incorporating the measurement information.

Typically, a learned reconstruction method would be trained for a single measurement operator. However, for radio interferometry, due to varying visibility coverage, the measurement operator changes with each observation. The challenge lies in creating a model that can reconstruct images for the **true visibility coverage** of an observation, while minimising the training time. In order to do this we propose the following strategies:

1. **True coverage**: model fully trained on true visibility coverage. Model needs to be fully retrained for every observation.
2. **Single coverage**: model trained using a single (different) visibility coverage. Model only needs to be trained once.
3. **Single coverage transfer**: model trained using a single (different) visibility coverage, then fine-tuned using transfer learning using the **true coverage** for every observation.
4. **Distribution of coverages**: model trained using a set of a 100 random visibility coverages. Model only needs to be trained once.
5. **Distribution of coverages transfer**: model trained using a set of a 100 random visibility coverages, then fine-tuned using transfer learning using the **true coverage** for every observation.

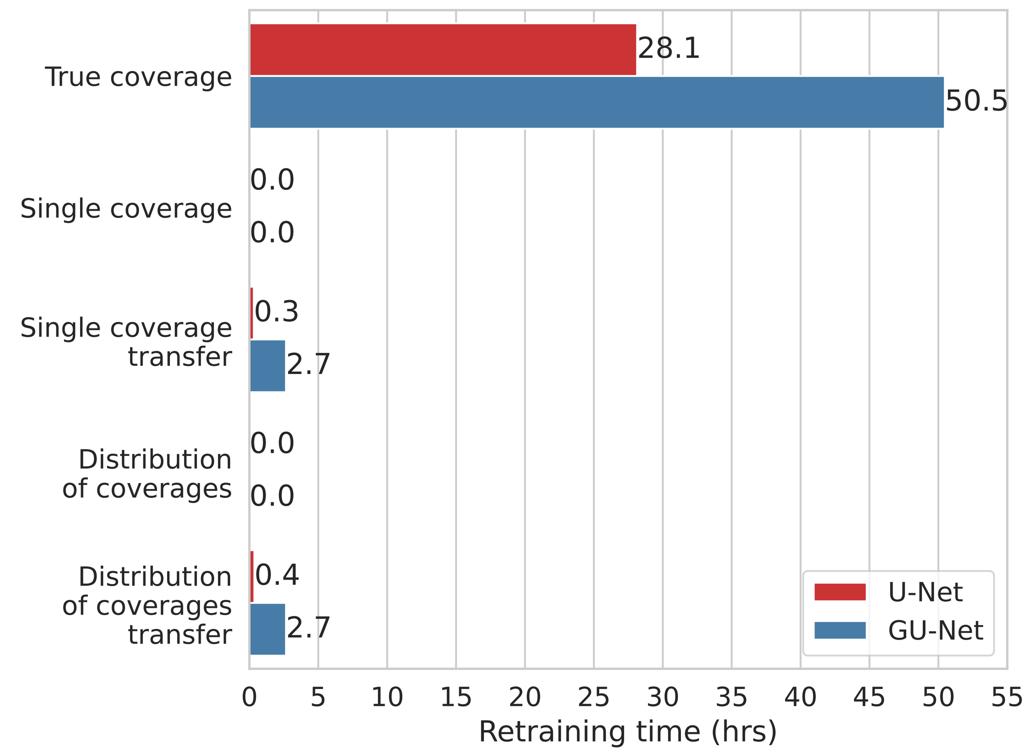

Figure: The time needed to train a network for the proposed strategies (including data augmentation).

4. Experiment

The networks are trained using 2000 galaxy images generated from the **IllustrisTNG simulations** [4]. The datasets are augmented using random flips, rotations and Gaussian random noise corresponding to an ISNR of 30dB. Networks are trained for 1000 epochs, with the ADAM optimizer, and a batch size of 20.

The models that are fine-tuned using transfer learning are trained for 100 epochs using 2000 non-augmented training images.

For evaluation, the models are (if necessary) rebuild with the measurement operator with the **true visibility coverage** and applied to measurements from the train (2000 images) and test (500 images) sets. The trained models are also applied to out-of-sample data not considered in training (30 Doradus).

5. Results: Reconstruction quality

Using the true visibility coverage will obviously lead to highest reconstruction fidelity but requires full retraining for each observation.
For the GU-Net in particular, we reach close to the benchmark reconstruction fidelity of the true visibility coverage but with little to no further training.

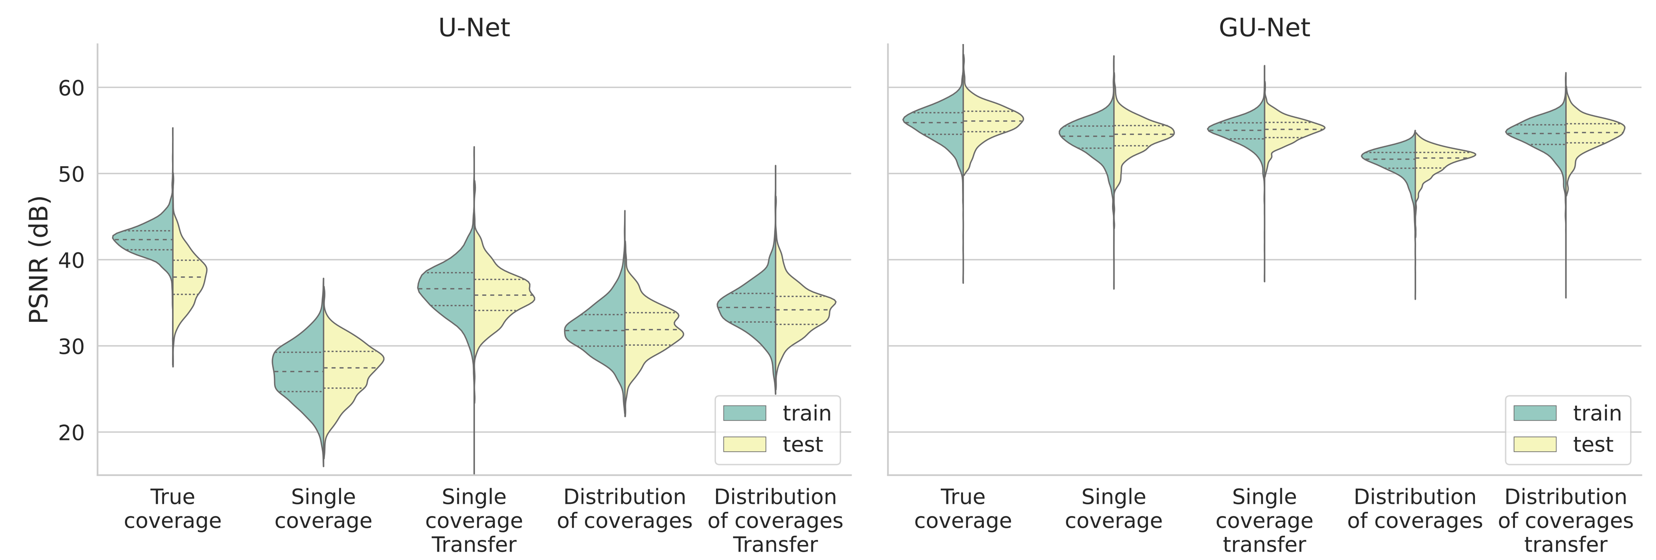

Figure: Reconstruction quality measured in peak signal-to-noise ratio (PSNR) for the different training strategies for both the U-Net and the GU-Net architectures.

6. Results: Example reconstruction of image from the IllustrisTNG simulation

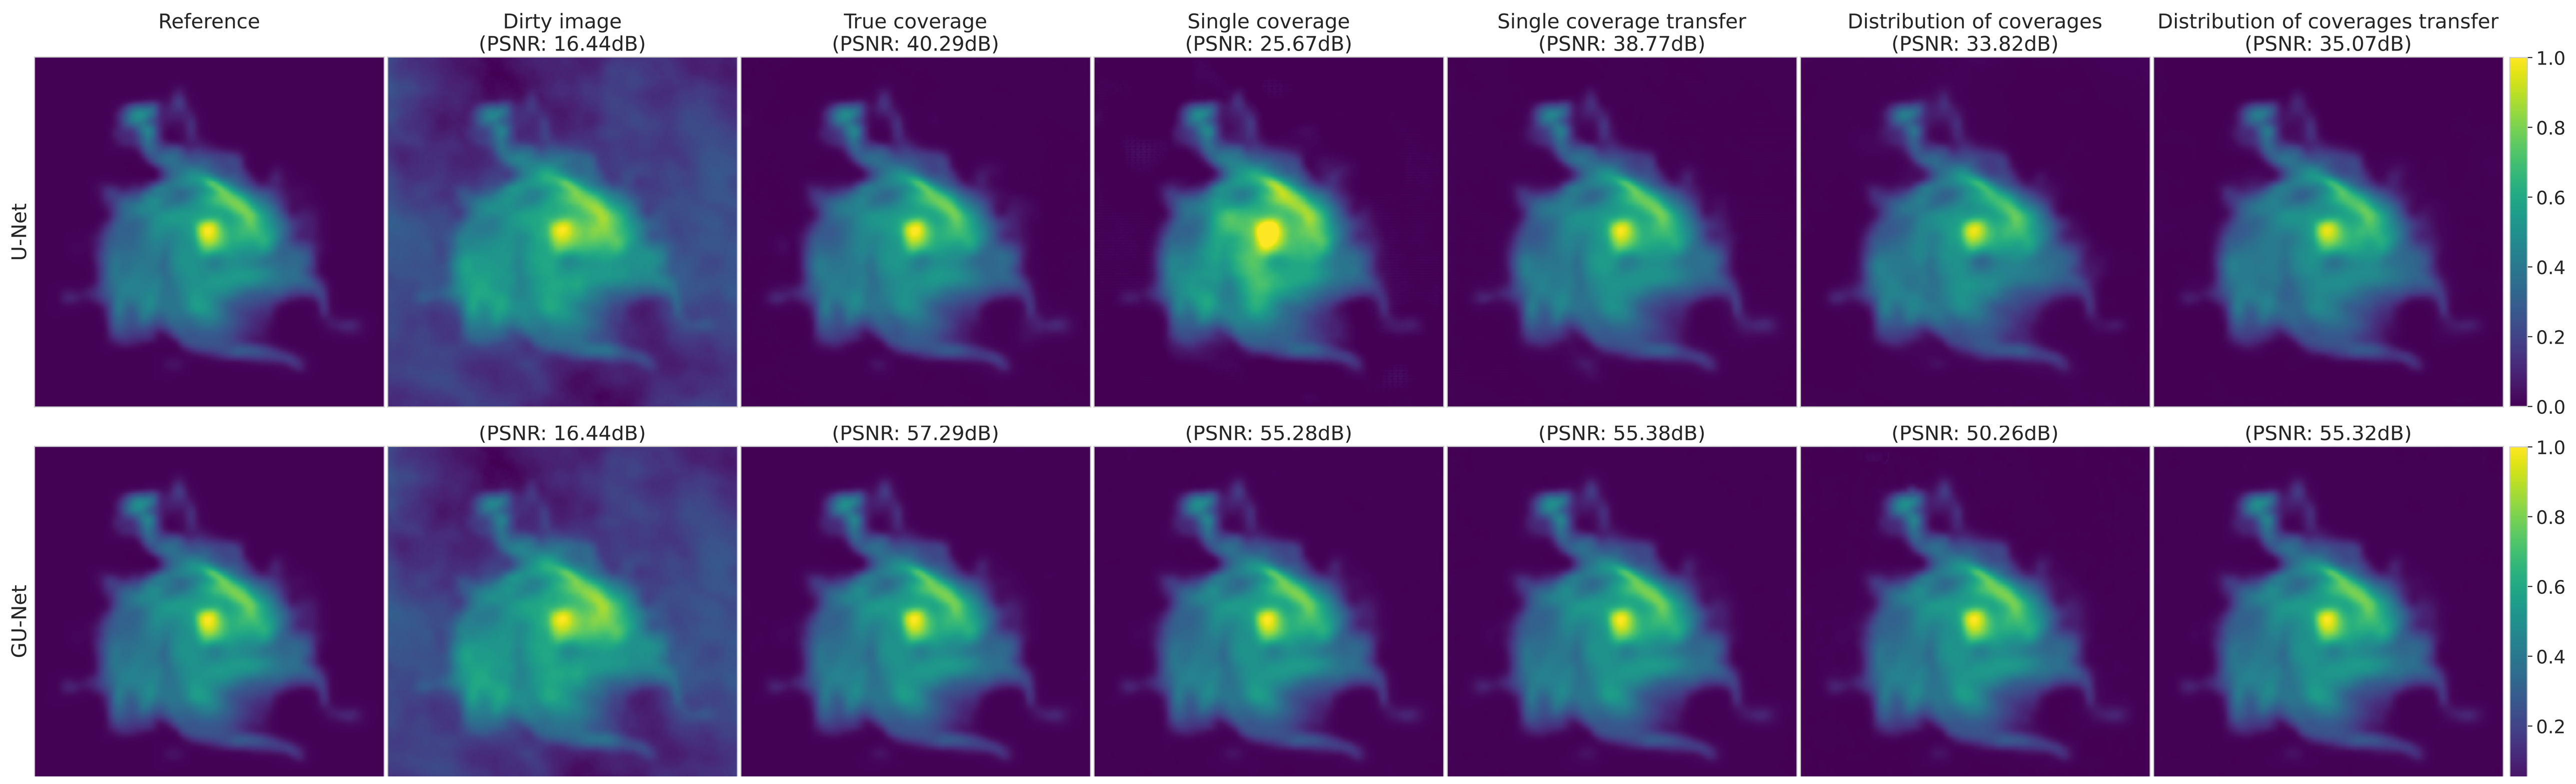

Figure: Reconstructions using the trained variants of the U-Net *(top)* and GU-Net *(bottom)* from simulated measurements of galaxy images created from the IllustrisTNG simulations

7. Results: Out-of-sample reconstruction of 30 Doradus

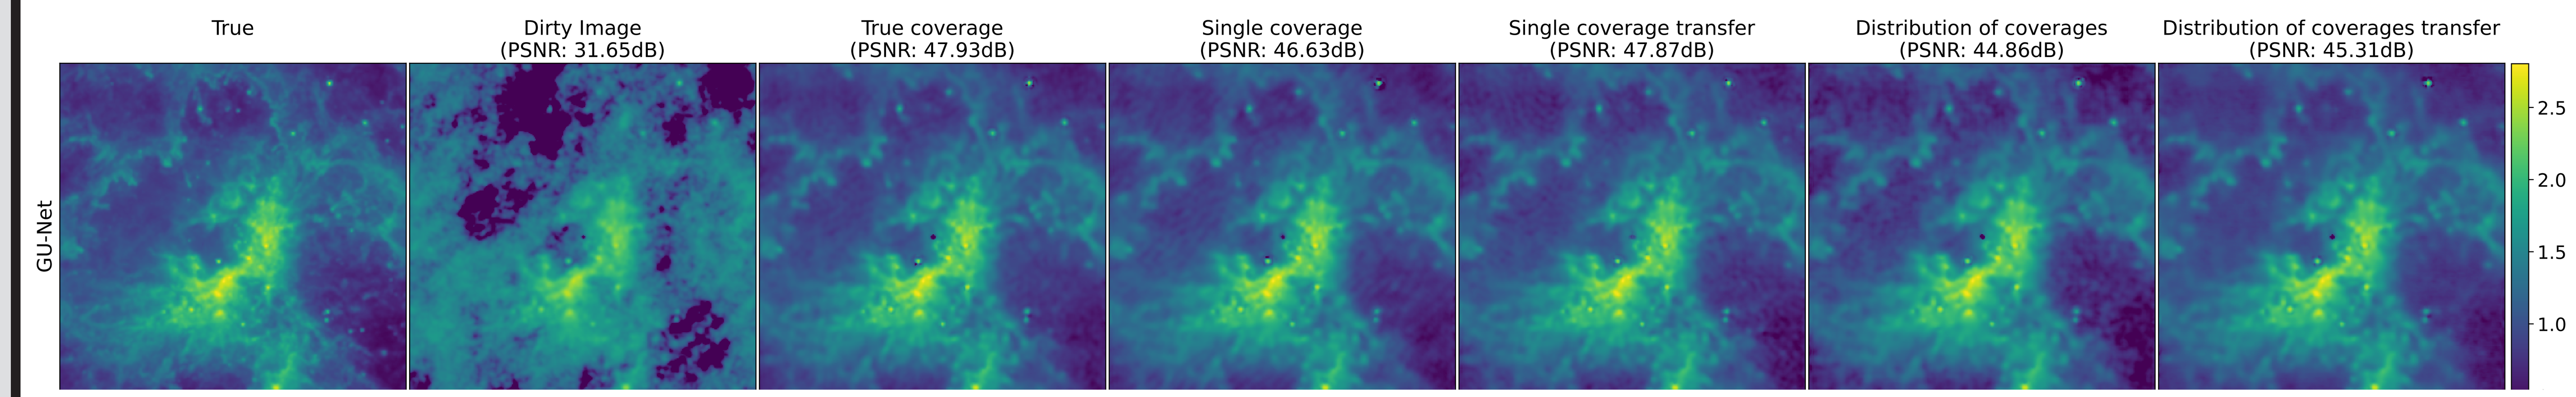

Figure: Reconstructions using the trained variants of the GU-Net from simulated measurements of an out-of-sample image of 30 Doradus, displayed on a log-scale.

References

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

[2] Dutt & Rohklin 1993, "Fast Fourier Transforms for Nonequispaced Data"

[3] Ronneberger et al. 2015, "U-Net: Convolutional Networks for Biomedical Image Segmentation"

[4] Nelson et al. 2019, "The IllustrisTNG Simulations: Public Data Release"

E-mail: matthijs.mars.20@ucl.ac.uk