

Learned radio interferometric imaging for varying visibility coverage

Matthijs Mars,^{1*} Marta M. Betcke,² Jason D. McEwen^{1†}

¹*Mullard Space Science Laboratory (MSSL), University College London (UCL), Dorking RH5 6NT, UK*

²*Department of Computer Science, University College London (UCL), London WC1E 6BT, UK*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

With the next generation of interferometric telescopes, such as the Square Kilometre Array (SKA), the need for highly efficient reconstruction techniques to process the huge volumes of data that will be acquired is particularly acute. Recent advances have shown that data-driven imaging approaches can simultaneously increase reconstruction quality while significantly reducing computational cost compared to alternative approaches. The challenge in designing these learned data-driven approaches for radio interferometry comes from the fact that the visibility coverage of the telescope is different for every observation as it depends on the telescope configuration, pointing direction, and duration of the observation. Learned methods must typically be retrained for each specific observation to account for these variations in visibility coverage, amounting to a large computational overhead. In this work we develop learned post-processing and unrolled iterative methods that support varying visibility coverages, proposing training strategies to make these methods agnostic to variations in visibility coverage with minimal to no retraining by, e.g., training on a distribution of coverages or by using transfer learning to fine-tune the network to the specific coverage. Learned post-processing techniques are heavily dependent on the prior information encoded in training data and generalise poorly to other coverages. In contrast, unrolled iterative methods, which include the telescope measurement operator inside the network, achieve state-of-the-art reconstruction quality and computation time, generalising well to other coverages and requiring little to no fine-tuning. Furthermore, the suitably trained methods generalise well to more realistic radio observations and are able to reconstruct the high dynamic range of these images.

Key words: machine learning – image processing – interferometric imaging

1 INTRODUCTION

To study radio sources in the Universe at high-resolution there is a need for radio telescopes with a very large aperture that cannot be achieved with a single dish. Instead, radio interferometry forms a telescope using multiple dishes or antennae spaced out over long distances to synthesise an aperture large enough to image the Universe in enough detail to uncover its mysteries. By combining the measurements from two of these antennae, which is referred to as a baseline, a so-called visibility can be measured, corresponding to one Fourier component of the sky. The inversion of this measurement process is an ill-posed problem as the radio interferometric telescope samples the Fourier domain non-uniformly, with a finite number of samples, and are contaminated by measurement noise. By projecting the acquired data to the image domain one recovers what is called a dirty image, a reconstruction of the target convolved with the point-spread function (PSF) of the telescope. Typically, various algorithms are used to reconstruct the image from acquired visibilities to yield a more accurate clean image of the true radio sky (e.g. CLEAN, Högbom 1974).

However, processing the vast volume of data generated by current large radio interferometers, such as the LOW-Frequency ARray

(LOFAR, Van Haarlem et al. 2013), the Murchison Widefield Array (MWA, Tingay et al. 2013), and the Australian SKA Pathfinder (ASKAP, Hotan et al. 2021), already presents immense computational challenges. The next generation of interferometric telescopes, such as the Square Kilometre Array (SKA, Dewdney et al. 2009), promises to revolutionise radio astronomy by offering unparalleled sensitivity, resolution, and sky coverage across a large frequency range. With the even larger volume of data generated by the SKA, yielding terabytes of data per second, there is an ever rising need for image reconstruction techniques that are able to process this huge amount of data efficiently.

Current image reconstruction techniques such as CLEAN (Högbom 1974) and its multi-scale variants (Bhatnagar & Cornwell 2004; Bhatnagar et al. 2008; Stewart et al. 2011) can be computationally costly and also require supervision, provide suboptimal reconstruction quality, and cannot provide any uncertainty quantification.

In contrast, techniques based on compressive sensing have demonstrated high quality image reconstructions providing both sharp and smooth image features (e.g. Carrillo et al. 2012; Pratley et al. 2018). Highly distributed algorithms have been developed for these approaches to efficiently reconstruct images for large numbers of visibilities (e.g. Pratley et al. 2019a) as well as to quantify the uncertainties (Cai et al. 2018a,b). Nonetheless, these methods need multiple applications of the measurement operator and its adjoint, which amounts to a large computational cost.

* E-mail: matthijs.mars97@gmail.com

† E-mail: jason.mcewen@ucl.ac.uk

Recently, learned approaches have gained in popularity, using training data sets as prior information for reconstructions. One particular approach of these learned methods implements learned regularisation, replacing the conventional regularisers in optimisation methods with learned priors. The most popular and promising method is to use so-called Plug-and-Play (PnP, Venkatakrishnan et al. 2013) denoisers that can be trained on different data domains, foregoing the need for large domain-specific data sets. Since these methods retain the decoupling of the prior and measurement operator they are robust to changes in the measurement operator and measurement noise level. In particular, PnP priors have shown promising results in radio imaging (Terris et al. 2023; Liaudat et al. 2023; Terris et al. 2022; Dabbech et al. 2022; Wilber et al. 2023). Furthermore, recent work has demonstrated how PnP methods can be extended to provide efficient uncertainty quantification (Liaudat et al. 2023). However, these methods require iterative reconstruction and so remain relatively computationally demanding.

Alternatively, one can use post-processing methods that use a denoising network to post-process an image created by using an approximate inverse of the measurement operator. In radio imaging this initial image is typically the dirty image and post-processing has been considered in several works (e.g. Allam Jr 2016; Terris et al. 2019; Gheller & Vazza 2021; Connor et al. 2022; Mars et al. 2023). Such post-processing networks depend more heavily on a correct model of the inverse of the measurement process and are more reliant on the prior information in the training set being representative of the actual observations. However, these methods are much faster than the learned regularisation methods as they only require a single evaluation of the measurement operator.

Learned unrolled iterative methods typically closely resemble iterative optimisation techniques, while reducing the number of iterations with the measurement operator, have shown promise in both reconstruction performance and in robustness to changes in the measurement operator. These methods have also produced promising results in astronomical imaging (e.g. Mars et al. 2023) and strike a balance between speed and quality of reconstruction. Inclusion of the measurement operator in the network architecture allows the network to retain some robustness to changes and uncertainty in the measurement operator (Boink et al. 2019, 2020; Maier et al. 2019), however results are conditioned on the measurement operator used during training.

One essential challenge in radio interferometric imaging is the variability of the Fourier domain coverage of telescopes, which is referred to as the uv -coverage. This is due to the coverage of the telescope depending on factors such as the pointing of the telescope and the rotation of the earth. This variability in the uv -coverage inevitably leads to a varying PSF for each observation and can therefore have a large impact on the quality of the reconstruction. Since the measurement operator of the telescope changes for every observation, any approach reliant on the measurement operator in training would typically have to be retrained for each uv -coverage. While robustness to modest variations in the PSF for these methods has been investigated, e.g., by applying transformations to the PSF that model observational effects (Connor et al. 2022), the use of post-processing and particularly learned unrolled iterative methods on different uv -coverages is still outstanding. Here we attempt to fill in this gap and study reconstruction methods that require minimal to no retraining for each observation by, e.g., training on a distribution of uv -coverages or by using transfer learning to adapt the network to the specific uv -coverage.

In this paper we discuss different strategies for training these learned reconstruction methods to determine how easily adapted they

are to the varying uv -coverages. We compare methods that require no further training as well as methods that apply transfer learning with additional fine-tuning with only use a small amount of retraining, compared to completely retraining the network for every different observation. We show that with only a small amount of overhead we can fine-tune these methods to the actual observations' uv -coverages, leading to a significant improvement in reconstruction quality.

The remainder of this paper is structured as follows. In Section 2 we discuss the interferometric imaging problem and how the measurement process can be modelled. Section 3 presents the different reconstruction approaches and the various training strategies that we develop for these learned methods. In Section 4 we demonstrate the performance of the different training strategies on the reconstruction of our simulated data set as well as simulated measurements from an actual radio observation. We discuss the implications of our results and the difficulties in applying these methods to real observations as well as how to potentially address these problems. Our findings are summarised in Section 5.

2 RADIO INTERFEROMETRIC IMAGING

The monochromatic interferometric measurement process for a radio telescope is described as (see, e.g., Thompson et al. 2017)

$$\mathcal{V}(u, v, w) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{\sqrt{1-l^2-m^2}} A_N(l, m) I(l, m) \times \exp \left\{ -i2\pi \left[ul + vm + w \left(\sqrt{1-l^2-m^2} - 1 \right) \right] \right\} dl dm, \quad (1)$$

where $\mathcal{V}(u, v, w)$ is the complex visibility, $A_N(l, m)$ is the primary beam pattern, and $I(l, m)$ is the sky brightness distribution. The coordinates (u, v, w) depend on the coordinates of the baselines and the coordinates (l, m) are directional cosines on the celestial sphere. For telescopes with a small field-of-view the baselines are all approximately co-planar (i.e. $w = 0$) and the measurement process reduces to the 2D Fourier transform

$$\mathcal{V}(u, v, 0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{A_N(l, m) I(l, m)}{\sqrt{1-l^2-m^2}} e^{-i2\pi(ul+vm)} dl dm. \quad (2)$$

However for radio telescopes with a large field-of-view, the baselines are not co-planar and the measurement process is typically approximated by adding operations that account for the curvature by adding w -projection or w -stacking terms (e.g. Cornwell et al. 2008; Tasse et al. 2013; Offringa et al. 2014; Pratley et al. 2019b). In this work we limit ourselves to the co-planar case however the methods presented can be extended (at an increased computational cost) by adding these additional corrections.

2.1 Interferometric measurement model

The discretised version of Equation 2 can be modelled by a so-called non-uniform discrete Fourier transform (NUDFT) mapping from N pixels in the observed image to M non-uniformly distributed visibilities, with $\Phi : \mathbb{R}^N \rightarrow \mathbb{C}^M$ (Pratley et al. 2018). The NUDFT can be approximated by the non-uniform fast Fourier transform (NUFFT, Duijndam & Schonewille 1997; Jackson et al. 1991; Fessler & Sutton 2003) which uses a gridding operation to interpolate the visibilities onto a regular grid and then uses the fast Fourier transform (FFT) to calculate the Fourier transform of the image. The forward NUFFT operator can be expressed as a series of operators

$$\Phi = \mathbf{GFZD}, \quad (3)$$

where $\mathbf{D} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is the diagonal matrix that contains a correction in the image domain for the gridding operation in the Fourier domain, $\mathbf{Z} : \mathbb{R}^N \rightarrow \mathbb{C}^{\alpha^2 N}$ a zero-padding operation that pads and therefore enlarges the image by a factor of α in each direction to upsample the Fourier domain, $\mathbf{F} : \mathbb{C}^{\alpha^2 N} \rightarrow \mathbb{C}^{\alpha^2 N}$ is the unitary FFT, and $\mathbf{G} : \mathbb{C}^{\alpha^2 N} \rightarrow \mathbb{C}^M$ the degriding operation that uses an interpolation kernel to interpolate the visibilities from the regular grid to the non-uniform measurements.

The adjoint of the measurement operator, $\Phi^* : \mathbb{C}^M \rightarrow \mathbb{R}^N$, can be expressed by using the adjoint of the individual operators in Equation 3 as

$$\Phi^* = \mathbf{D}\mathbf{Z}^* \mathbf{F}^\dagger \mathbf{G}^*, \quad (4)$$

with $\mathbf{G}^* : \mathbb{C}^M \rightarrow \mathbb{C}^{\alpha^2 N}$ the gridding operation that uses a gridding kernel to grid the measurements to the uniform grid, $\mathbf{F}^\dagger : \mathbb{C}^{\alpha^2 N} \rightarrow \mathbb{C}^{\alpha^2 N}$ the inverse unitary FFT, and $\mathbf{Z}^* : \mathbb{R}^{\alpha^2 N} \rightarrow \mathbb{R}^N$ the adjoint of the zero-padding operation, cropping the image back to its original size, and $\mathbf{D}^* = \mathbf{D}$ because of the self-adjointness of the real diagonal matrix correcting for the gridding operation.

To obtain an approximate inverse of the measurement process, we include measurement weights in the adjoint operation. In radio interferometry the measurements are typically weighted using either natural, uniform or robust weighting schemes (Taylor et al. 1999). Natural weights corresponds to the inverse of the standard deviation of the measurement uncertainty $\mathbf{W}_{k,k}^{\text{Natural}} = \sigma_k^{-2}$, with σ_k the standard deviation on the measured visibility $\mathcal{V}(u_k, v_k, 0)$ for the k -th visibility, and is often also referred to as whitening. Uniform weighting compensates for the increased sampling density at lower spatial frequencies by additionally weighing the measurements by the density of the sampling distribution, which causes the reconstructed image to be dominated by low spatial frequency power; $\mathbf{W}_{k,k}^{\text{Uniform}} = \sigma_k^{-2} N_s^{-1}(k)$, with $N_s(k)$ the number of measurements in a region around the measurement $\mathcal{V}(u_k, v_k, 0)$. Lastly, robust weighting, also referred to as Briggs weighting, considers a trade-off between natural and uniform weighting, controlled by a robustness parameter (Briggs 1995). The pseudo-inverse of the measurement operator applied to the measurements is typically referred to as the dirty image, and is calculated using

$$\Phi^\dagger = \mathbf{D}\mathbf{Z}^* \mathbf{F}^\dagger \mathbf{G}^* \mathbf{W}. \quad (5)$$

A more detailed explanation of the measurement process and the NUFFT operator can be found in, e.g., Pratley et al. (2019a); Mars et al. (2023).

2.2 Inverse problem

The radio interferometric measurement formation is given as

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{n}, \quad (6)$$

with $\Phi : \mathbb{R}^N \rightarrow \mathbb{C}^M$ the measurement operator mapping from the true image $\mathbf{x} \in \mathbb{R}^N$ to the non-uniformly distributed measurements $\mathbf{y} \in \mathbb{C}^M$ with $\mathbf{n} \in \mathbb{C}^M$ the measurement noise.

The CLEAN algorithm solves the corresponding inverse problem using a greedy algorithm that aims to iteratively remove the brightest point sources from the observation in order to build up a clean sky model (Högbom 1974). Many variations of the algorithm exist that aim to improve or provide reconstruction of resolved, extended or polarised sources as well as multi-frequency synthesis (Clark 1980; Schwab 1984; Steer et al. 1984; Sault et al. 1996; Cornwell 2008; Offringa et al. 2014; Pratley & Johnston-Hollitt 2016). These variants can be used to reconstruct images with a high dynamic range and

are widely used in radio interferometry. While the algorithm is very versatile, it has many parameters that need to be tuned for each observation and reconstructing using CLEAN is by some considered more an art than a science.

Popular alternative inversion approaches involves the use of variational regularisation in order to reconstruct the image by minimising

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \mathcal{L}(\Phi \mathbf{x}, \mathbf{y}) + \lambda \mathcal{S}(\mathbf{x}), \quad (7)$$

which provides a trade off between the data fidelity term $\mathcal{L}(\Phi \mathbf{x}, \mathbf{y})$ and the regularisation term $\mathcal{S}(\mathbf{x})$, which encodes prior information of the images and is weighted by the regularisation parameter λ . In radio interferometry, a typical data fidelity term is the squared ℓ_2 -norm, i.e. $\mathcal{L}(\Phi \mathbf{x}, \mathbf{y}) = \|\Phi \mathbf{x} - \mathbf{y}\|_{\ell_2}^2$, and a widely used regularisation term is an ℓ_1 -norm on an overcomplete dictionary of wavelet bases, e.g. consisting of a Dirac basis and Debauchies wavelets bases (Daubechies 1992), as used in the sparsity averaging reweighted analysis (SARA, Carrillo et al. 2012) algorithm. Alternatively, this problem can be formulated in the constrained setting

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \mathcal{S}(\mathbf{x}) \quad \text{s.t.} \quad \|\Phi \mathbf{x} - \mathbf{y}\|_{\ell_2}^2 \leq \epsilon. \quad (8)$$

The optimisation problems in Equations 7 and 8 are typically solved using proximal optimisation techniques, such as the fast iterative shrinkage-thresholding algorithm (FISTA, Beck & Teboulle 2009), the alternating direction method of multipliers (ADMM, Boyd 2010), or the primal-dual hybrid gradient method (PDHG, Chambolle et al. 2018) and a variety of these methods have been applied to radio interferometry (e.g. Pratley et al. 2018, 2019a).

Learned reconstruction approaches aim to replace some parts of these traditional optimisation techniques with neural networks in order to incorporate a data-driven prior and/or to speed up the reconstruction process. Learned regularisation approaches use neural networks to capture the prior information in the variational problem. In particular, PnP methods replace the proximal problem formulation in proximal splitting algorithms with a denoiser (Venkatkrishnan et al. 2013; Ryu et al. 2019). These denoisers are typically trained on data sets outside of the target domain, e.g. natural images, and encode more general prior information about the denoising of images. The iterative updates based on both these more general priors and the data fidelity term in the optimisation problem allow these methods to generalise well to unseen data. While these methods are able to achieve high reconstruction quality they are computationally expensive as they require many iterations and therefore evaluations of the measurement operator to converge to a solution.

Instead, learned post-processing methods (e.g. Chen et al. 2017; Jin et al. 2017; Yi & Babyn 2018) aim to drastically speed up the reconstruction process by training a neural network to post-process an initial reconstruction. In radio interferometry, this initial reconstruction is typically the dirty image obtained by the pseudo-inverse of the measurements. Figure 1 shows the typical pipeline of a learned post-processing method in radio interferometry. Because the network is decoupled from the measurement process, these methods can be trained using an image-to-image network making them much faster to train and evaluate. However, these methods are heavily reliant on the quality of the initial reconstruction and dependant on the artefact distribution of the uv -coverage(s) used for the training data.

Learned unrolled iterative methods combine the data-driven priors with model based updates in order to provide a balance between the speed of learned post-processing methods and the quality of learned regularisation methods. They typically mimic proximal optimisation algorithms by unrolling a small number of iterations of the




Figure 1. The learned post-processing approach creates a initial reconstruction from the measurements using a model of the telescope. The resulting dirty image is passed through a post-processing neural network to create the reconstruction. Our U-Net model takes this approach.

iterative process and replacing some of the proximal operators with neural networks (Gregor & LeCun 2010). Figure 2 shows a diagram summarising these methods. While a large portion of these methods closely follow the traditional optimisation methods (e.g. Adler & Öktem 2018), some methods have been proposed that incorporate the measurement process more directly into the network. For instance, Hauptmann et al. (2020) proposed a multi-scale learned iterative method that includes the measurement information at different resolution scales in the multi-scale U-Net architecture (Ronneberger et al. 2015). Similarly there are methods that see the inclusion of the gradient of a data fidelity term added into this U-Net structure at several resolution scales (Trent 2020; Mars et al. 2023). The benefit in these methods is the reduced computation and memory overhead as at the lower scales of the network the measurement operator is evaluated at a lower resolution and therefore more efficiently.

3 LEARNED RECONSTRUCTION TECHNIQUES

In order to provide fast image reconstruction of radio interferometric observations and to exploit the enhanced expressiveness of learned data-driven priors, we propose to use either learned post-processing or learned unrolled iterative methods. These approaches minimise the number of applications of the measurement operator and its adjoint, which is the dominating factor in the computational cost of the reconstruction. In particular, we consider the approaches introduced in our prior work (Mars et al. 2023), where these learned methods were used to provide efficient reconstruction for an optical interferometric telescope called SPIDER. While in that case the uv -coverage was fixed we now discuss how these methods can be adapted to work with varying uv -coverages. In order to do this we first discuss the reconstruction models that we use in this article and then propose a variety of different strategies to train the networks such that they generalise to any uv -coverage. We expect the unrolled iterative method to be more robust to changes in the uv -coverage as it includes the measurement operator in the network architecture, while the post-processing method is more reliant on the prior information in the training data being representative of the actual observations.

3.1 Learned post-processing method

The first reconstruction approach we consider is a learned post-processing approach that takes the dirty reconstruction as input and uses a U-Net architecture to remove artifacts from this initial reconstruction. A schematic representation of the network architecture can be found in Figure 3. The reconstruction can be formulated as:

$$\mathbf{x}^\star = \Phi_\theta^\dagger \mathbf{y} = \Lambda_\theta \Phi^\dagger \mathbf{y}, \quad (9)$$




Figure 2. The learned unrolled iterative approach also creates an initial reconstruction using the telescope model and passes this to the reconstruction network. However, it also uses the measurements and the telescope model in conjunction with the neural network to update the image and provide the final reconstruction. Our GU-Net model takes this approach.

where $\Lambda_\theta : \mathbb{R}^N \rightarrow \mathbb{R}^N$ denotes the reconstruction network and $\Phi^\dagger : \mathbb{C}^M \rightarrow \mathbb{R}^N$ the pseudo-inverse of the measurement operator using natural weighting. This post-processing is computationally efficient since the measurement operator is evaluated only once to create the initial reconstruction. However, the final reconstruction quality is limited by the quality of this initial reconstruction. Besides this, these reconstruction networks will learn how to remove errors that correspond specifically to the uv -coverage used in the training data. The implementation of the U-Net used in this article is similar to that of the U-Net used in Mars et al. (2023), with the addition of a residual connection between the input and output of the network. The initial reconstructions are created using the appropriate measurement operator for the uv -coverage of the observation.

3.2 Learned unrolled iterative method

The second approach we consider is a learned unrolled iterative method that uses the measurement operator and measurements of the observation in the neural network to improve the reconstruction quality. We extend the U-Net architecture by introducing measurement information at each scale of the network. To achieve this we use so-called sub-scale operators $\Phi_i : \mathbb{R}^{N_i} \rightarrow \mathbb{C}_i^{M_i}$ that map from a reduced image space $\mathbf{x}_i \in \mathbb{R}^{N_i}$ to low-pass filtered visibilities $\mathbf{y}_i \in \mathbb{C}^{M_i}$. The use of sub-scale operators is discussed in more detail in Mars et al. (2023). These sub-scale operators are used to include the gradient of the data fidelity term $\mathcal{L}(\Phi_i \mathbf{x}_i, \mathbf{y}_i) = \frac{1}{2} \|\Phi_i \mathbf{x}_i - \mathbf{y}_i\|_{\ell_2}^2$ in the multi-scale U-Net:

$$\nabla_{\mathbf{x}_i} \mathcal{L}(\Phi_i \mathbf{x}_i, \mathbf{y}_i) \propto \Phi_i^* (\Phi_i \mathbf{x}_i - \mathbf{y}_i), \quad (10)$$

where $\mathbf{x}_i \in \mathbb{R}^{N_i}$ is the current iterate at scale i of the network and $\mathbf{y}_i \in \mathbb{C}^{M_i}$ are the low-pass filtered visibilities at scale i of the network.

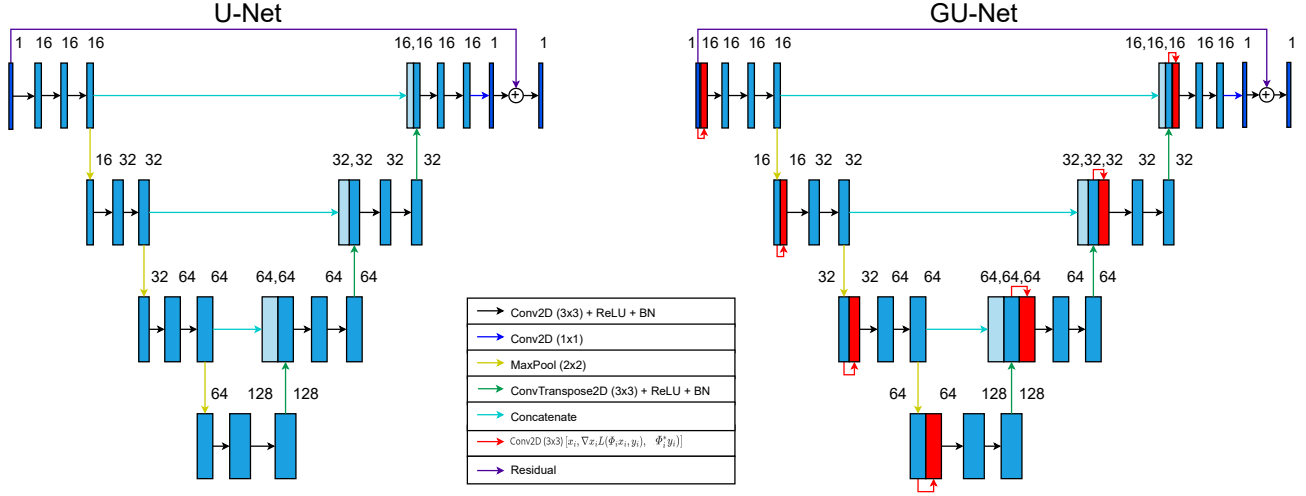


Figure 3. Left: The U-Net architecture for the learned post-processing approach, with the input of the model being the dirty image. Right: The GU-Net architecture for the learned unrolled iterative approach that incorporates both the measurements and the measurement operator in the network through the (sub-scale) gradients. The model-based update takes the image in the first channel and calculates its gradient (Equation 10) and the sub-sampled dirty image (Equation 11), combines them using a convolutional layer, and concatenates this with the original channels.

In addition to the gradient information we also include the sub-sampled dirty image, $\mathbf{x}_{i,\text{dirty}} \in \mathbb{R}^{N_i}$ at each of the gradient update steps

$$\mathbf{x}_{i,\text{dirty}} = \mathbf{\Phi}_i^\dagger \mathbf{y}_i. \quad (11)$$

We call the resulting network a Gradient U-Net (GU-Net, Mars et al. 2023). While this network includes several gradient updates (in this case seven), only two of these are at the full image scale and the rest are at the sub-scales of the network, which can be computed at a much lower computational cost because of the reduction in image size and restriction of the Fourier space.

In contrast to our previous work (Mars et al. 2023), we do not include a filtered gradient, corresponding to the gradient of a weighted data fidelity term, in the network architecture as its contribution is marginal compared to the increase in computational time and thus training time that it requires. We also include a residual connection between the input and output of the network. The GU-Net architecture is displayed in Figure 3.

3.3 Training strategies

In order to train a reconstruction network that is robust against the variation in the uv -coverage of the radio telescope, we propose several different training strategies. In these strategies we differentiate between the *true coverage*, which is the uv -coverage of the observation we are trying to reconstruct, and other coverages drawn from a similar distribution as the true coverage, yet different to the actual observation.

The uv -coverages are drawn from a Gaussian distribution, yet the training strategies could work with coverages drawn from a distribution corresponding to a specific telescope in the future. We also discuss the results of fine-tuning a model to a realistic uv -coverage from the MeerKAT telescope using transfer learning in Section 4.4.

3.3.1 Training on the true coverage

The first training strategy is to train networks on the true coverage of the observation. The benefit of this strategy is that it is conditioned on the exact uv -coverage of the observation, yet this can only be known after the observation. Therefore, the networks would have to be fully trained after each observation. This strategy can be seen as the "oracle" strategy as it is the best case scenario if you have perfect knowledge of the observation. We will use this strategy as a best case benchmark to evaluate the performance of the other strategies.

3.3.2 Training on a single coverage

Instead of training on the true coverage, we can also use training data that is generated with a single alternate coverage that is drawn from a distribution of coverages representative of the radio telescope. At time of actual reconstruction, the measurement operators that were used during training can be swapped out for operators that use the true coverage of the observation. This approach is expected to perform better using reconstruction approaches that rely less on the prior information related to the measurement model in the training data and that leverage the measurement operator in the reconstruction process, such as the GU-Net approach.

3.3.3 Training on a distribution of coverages

Alternatively, we can learn a reconstruction approach that is conditioned on the distribution of coverages representative of the radio telescope. In order to do so, we create training data using a different random uv -coverage for every batch to train our networks. The performance of this method depends on how well the true coverage is characterised by the distribution of coverages used for training. If there is a large variance in the distribution of coverages, the model will have a harder time learning the underlying connection between the measurements and the image.

Training a post-processing method using this strategy is trivial as

the network can be trained as an image-to-image network without explicit calls to the measurement operator in the training process. However, for the learned unrolled iterative method, the measurement operator is explicitly included in the training process, posing a challenge as it needs to be varied throughout training. In order to achieve this there are three possible approaches. The first option is to swap the measurement operator in the model during training, which is not supported by most deep learning frameworks and therefore requires recompiling the execution graph whenever the operator needs to be swapped, resulting in significantly increased training times. For the second option, the uv -coverage can be added as an input to the network in order to create the measurement operator as part of the training process, though this results in a large computational overhead in order to compute the coefficients for the (de)gridding of the non-uniformly distributed visibilities on the fly. For the third option, we can compile the network with all possible measurement operators and use a mask to select the correct operator for the current batch of measurements. In general, this results in memory overheads because of the large number of operators that need to be compiled into the network. For the case of the random Gaussian coverages, we use this third approach, generating a larger number of visibilities and randomly select a sub-selection of these for each batch of training. In this way, the memory overhead is reduced as the random visibilities will share some visibility measurements between them.

3.3.4 Transfer learning models

While the previous two strategies do not depend on the true coverage and therefore do not need any retraining at observation, we can use transfer learning to fine-tune the other strategies. This is achieved by using the weights of the models that were trained using the previous strategies as a starting point and fine-tuning these weights by training the model using a data set created using the true coverage. Since the number of training epochs needed to fine-tune the models is much smaller than the number needed to fully retrain the models, this results in a much more manageable training process. We can also forego the use of data augmentation as the risk of overfitting is smaller. The specifics of the data sets used for training and transfer learning as well as data augmentation are discussed in Section 3.4. This results in reconstruction networks adapted to the specific uv -coverage of the observation at a relatively modest cost of retraining for each observation.

3.4 Training and training data

To train the learned reconstruction networks we create a data set of measurements using images of simulated galaxies created from the IllustrisTNG simulations (Nelson et al. 2019b,a), which were obtained by binning and in-painting the image from the simulated particles as described in more detail in Mars et al. (2023). The images are augmented during training by randomly flipping and rotating them. We use an NUFFT operator to simulate radio interferometric measurements, as discussed in Section 2.1, from the 256×256 images, using an upsampling factor of 2 and Kaiser-Bessel kernels with a size of 6×6 pixels. The uv -coverage used is drawn from a 2D Gaussian distribution and results in a total of 32768 measurements. These measurements are contaminated with complex Gaussian noise with standard deviation

$$\sigma = \frac{\|\Phi\mathbf{x}\|_{\ell_2}}{\sqrt{M}} \cdot 10^{-\frac{\text{ISNR}}{20}}, \quad (12)$$

with the input signal-to-noise ratio (ISNR) of 30dB, and this noise is re-drawn for every epoch of training.

From the IllustrisTNG images, 2000 images are used for training and 1000 for testing. These image sets are sampled with the uv -coverages drawn from the same Gaussian distribution to create several different training sets. First, we create a train and test data set based on the true uv -coverage that will be used to train the "oracle" network and to evaluate the performance of the different models. Then, a second different uv -coverage is used to create a training data set for the single coverage model. Finally, a set of random uv -coverages is used to create a data set to train the distribution of coverages model. For this data set, every batch of training data is generated using a randomly generated coverage. The data set used for the transfer learning of the models uses the true coverage to create measurements, but does not use any type of data augmentation as considered for the previous data sets, such that the input-output pairs for the model only have to be generated once for the full training of these models, keeping the time needed for training low.

The networks are trained for 1000 epochs using the Adam optimizer (Kingma & Ba 2017) with a learning rate of 10^{-3} and a batch size of 20, using a mean squared error loss between the true image and the reconstructed image. Transfer learning using the non-augmented data set based on the true coverage is performed for 100 epochs. The data sets and the weights of the networks are 32-bit floating point numbers and the networks are trained on a single NVIDIA A100 GPU.

4 RESULTS OF NUMERICAL EXPERIMENTS

We evaluate the effectiveness of the proposed training strategies in terms of reconstruction quality, (re)training times, robustness to variation within the measurement distribution and ability to generalise to more realistic radio observations. First, we show that using these strategies reduces the time necessary to prepare a model for each observation as well as discussing the computational efficiency of the reconstruction methods. The quality of reconstructions is assessed on images from the IllustrisTNG data set and an example of a realistic radio observation in terms of the signal-to-noise ratio, $\text{SNR}(\mathbf{x}_{\text{true}}, \mathbf{x}_{\text{pred}}) = 20 \cdot \log_{10} \left(\frac{\|\mathbf{x}_{\text{true}}\|}{\|\mathbf{x}_{\text{pred}} - \mathbf{x}_{\text{true}}\|} \right)$ with respect to the true image. For the more realistic radio observations, we also evaluate the reconstructions on a logarithmic scale taking into account their extended dynamic range. For these evaluations we track the logarithmic SNR; $\log \text{SNR}(\mathbf{x}_{\text{true}}, \mathbf{x}_{\text{pred}}) = \text{SNR}(\log_{10}(\mathbf{x}_{\text{true}}), \log_{10}(\mathbf{x}_{\text{pred}}))$. Finally, we demonstrate how these models can be adapted for use with actual radio uv -coverages by transferring the learned models to a real telescope and evaluating their performance on a realistic radio observation.

4.1 Training and evaluation times

The main aim of the different training strategies is to reduce the time needed to (re)train a model for each observation, while still maintaining high reconstruction quality. Figure 4 shows the time needed to get a reconstruction network ready for image reconstruction. The training of the full models takes more than 1 day for the U-Net (on one A100 GPU with 40Gb VRAM) and more than 2 days for the GU-Net models and are therefore not feasible to be retrained for every observation. A large portion of this time (~ 24 hrs) comes from the creation and pre-augmentation of the data set using the specific uv -coverage(s) for training. Instead we have two strategies that require




Figure 4. The time necessary for training both the U-Net and GU-Net models (on one A100 GPU with 40Gb VRAM) once visibility coverage is available. The proposed approaches reduce post-observation retraining time from 28 and 51 hours to zero hours if not fine-tuned, for the U-Net and GU-Net respectively. If transfer learning is applied, fine-tuning the models, to improve reconstruction quality retraining time is reduced to 0.3 and 2.7 hours respectively. Note that the time for training the full model (the *true coverage* strategy) also contains the creation and pre-augmentation of the training data which took ~ 24 hrs.

no retraining: the single coverage and distribution of coverages. Furthermore, we have two approaches that only require a small amount of retraining in the form of transfer learning. These transfer learned models do not require a large and augmented data set to be created and can be trained in a matter of minutes or hours.

Table 1 shows the reconstruction time necessary for the U-Net and GU-net models once trained. The reconstruction is dominated by the evaluations of the measurement operator to create the initial reconstruction for both models and for the calculation of the gradients and dirty images in the GU-Net. Reconstruction time scales linearly with the number of full-scale evaluations of the operator. The computational cost of state-of-the-art iterative reconstruction methods also scales linearly with the number of evaluations of the measurement operator as these are evaluated at each iteration. However, these techniques typically require ~ 50 iterations or more to converge and are therefore much more computationally expensive than the proposed methods, which require only a few evaluations.

4.2 Reconstruction quality

We assess the reconstruction quality for the five training strategies by evaluating the SNR on reconstructions of images from the train and test set. Figure 5 shows the distribution of the reconstruction SNR for both the U-Net and the GU-Net on these simulated measurements.

The performance of the U-Net architecture is significantly worse than the GU-Net architecture, with the GU-Net architecture achieving a higher reconstruction quality for all the different training strategies, indicating that adding in measurement information in the reconstruction process significantly improves the reconstruction quality.

For the U-Net, we also notice that training the model using data from the true coverage results in a model that overfits to the training data set, as the reconstruction quality on the test set is significantly worse than the training set. While this could be improved by using a larger training set or a training set with more variance (or by in-

Table 1. The average reconstruction time by creating the dirty image using the pseudo-inverse, and the reconstructions using the U-Net and the GU-Net on an A100 GPU. We also detail the number of full-scale evaluations of the measurement operator and its adjoint.

Name	Operator evaluations	Reconstruction time (ms)
Pseudo-inverse	1	13.0 ± 0.4
U-Net	1	54.9 ± 1.6
GU-Net	7*	85.7 ± 1.5

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

cluding additional methods to combat overfitting), this tendency to overfit shows that this reconstruction approach itself relies heavily on the prior information in the training set and will not generalise well to images outside of the training distribution. Furthermore, we can see this in the performance of the U-Net trained on a single coverage, which generalises very poorly when used to reconstruct images from measurements using the true coverage. This is not unexpected since the U-Net reconstruction only includes the measurement operator explicitly in the initial reconstruction passed to the network. Therefore this method is biased towards the artefact distribution of the *uv*-coverage used in training. Training the U-Net using data of a distribution of coverages alleviates some of these issues as the network is able to learn the prior information for an arbitrary set of coverages with some performance loss compared to training on the actual coverage. A minimal amount of transfer learning increases the performance for both strategies, but the performance is still worse than fully retraining a model using data simulated using the true coverage. Longer retraining would likely further increase the performance of these models, but at some point we risk overfitting to the train set as we observed in the model trained on the true coverage.

The performance of the GU-Net is more uniform across the different strategies. With no to little retraining we achieve significantly closer results to training using the true coverage. The integration of the measurement operator inside the reconstruction network allows the network to learn the prior information irrespective of the coverage used. This is reflected in the performance of the single coverage model which generalises well when used to reconstruct images from measurements using the true coverage, without any retraining. The distribution of coverages strategy seems to generalise slightly worse to the true measurement sets, which is likely due to the fact that it is trying to learn a more general prior that works for any coverage. The discrepancy between these two strategies is directly linked to the distribution of *uv*-coverages used for training. In this case the errors induced by the coverage used for the single coverage network are similar enough to that of those induced by the true coverage resulting in good generalisation from one to the other. However, if the distribution of telescope *uv*-coverages has a significantly larger variance it would be expected that the single coverage model would not generalise as well to the true coverage and the distribution of coverages model might be preferred. Despite this, the performance of both models is improved using a small amount of transfer learning, bringing the reconstruction quality close to that of the oracle model trained using the true coverage.

Figure 6 shows example reconstructions using the U-Net for the different training strategies. While the reconstructions are visually similar the reconstruction quality as indicated by the SNR is significantly different and reflects the results in Figure 5. The difference in the reconstruction strategies can be seen in their accuracy of predicting the diffuse structures of the images as well as correctly

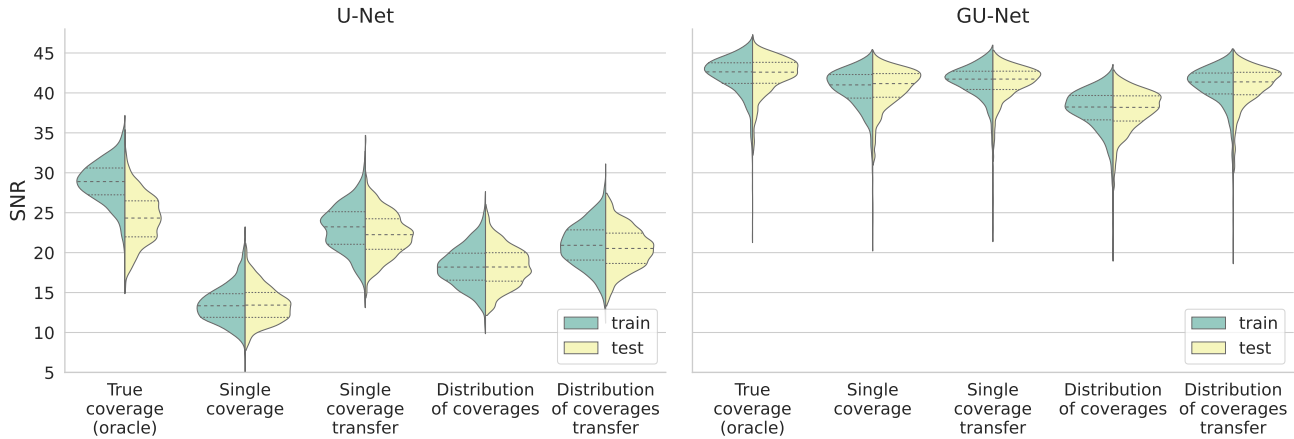


Figure 5. Reconstruction quality measured in SNR for the different training strategies for both the U-Net and the GU-Net architectures. The GU-Net models provide significantly higher reconstruction quality than the U-Net models and perform similarly across the different strategies. Without any retraining (the single and distribution of coverages strategies) these models lose only a small amount of reconstructive performance compared to fully retraining on the true coverage. Fine-tuning using transfer learning for a short time restores performance close to that of the true coverage. A similar trend across training strategies can be seen in the U-Net though the difference in performances is much larger and the transfer learning is not able to fully restore the reconstructive performance.

reconstructing the intensity of the brighter features. Besides that, we see that the model using the single coverage strategy produces reconstructions that contain hallucinations of diffuse structure that are not found in the ground truth (most notably in the fourth and fifth example reconstructions).

In Figure 7 we see the reconstructions using the GU-Net for the different training strategies. The different strategies produce visually similar looking reconstructions with high reconstruction quality. While the network using a distribution of coverages for training performs slightly worse than the other strategies, using a small amount of transfer learning improves the performance significantly to match that of the other strategies.

4.3 High dynamic range reconstruction quality

While the proposed methods are trained on images of simulated galaxies, the resulting networks should also be tested on more realistic radio observations. Therefore, to test the generalisation power of the different training strategies, we evaluate the performance of the models on an out-of-sample image of 30 Doradus. The dynamic range of this image is significantly larger than the simulated images and we therefore also evaluate the reconstructions on a logarithmic scale.

Figure 8 shows the reconstructions using the U-Net for the different training strategies. The models that are not trained or fine-tuned using data from the true coverage perform significantly worse for this post-processing approach, whereas we see that fine-tuning the models produces a higher quality of reconstructions. Nonetheless, the models are not able to reconstruct the dynamic range of the image correctly, with the reconstructions showing large amounts of noise and artefacts. This lack of performance traces back to the fact that the models rely heavily on the prior information in the training set and are not able to generalise well to new data with a higher dynamic range.

When looking at the reconstructions from the same measurements using the GU-Net in Figure 9, we see that the models are able to produce high quality reconstructions with a high dynamic range. Similar to on the simulated images, the methods not fine-tuned on data from the true coverage performs slightly worse but this difference

is easily mitigated by fine-tuning the models using transfer learning, yielding reconstructions that are of the same reconstruction quality as fully retraining the model. The explicit inclusion of the measurement operator in the reconstruction process allows the network to learn the prior information irrespective of the coverage used, which is why the GU-Net is able to generalise much better to out-of-distribution data as well as generalising to different uv -coverages. This generalisability makes it more suitable for application to real radio observations, since it is able to correct for differences between the training data set and the actual observations.

4.4 Transfer learning to MeerKAT uv -coverage

Finally, to investigate whether these reconstruction networks generalise from the Gaussian visibility coverages to realistic uv -coverages, we demonstrate that the same models can be used to reconstruct from simulated measurements using a uv -coverage from the MeerKAT telescope with 241920 visibilities as seen in Figure 10. In order to do this we fine-tune the GU-Net model trained on the single coverage by training it on a small data set created using the MeerKAT coverage. The model is trained for 50 epochs and the creation of the data took ~ 2 hours and the fine-tuning of the model took ~ 5 hours on an A100 GPU with 40GB VRAM. For this task we only consider the GU-Net architecture as the U-Net models do not generalise well to high dynamic range images as demonstrated in Section 4.3. The fine-tuned GU-Net is then used to reconstruct an out-of-sample image of 30 Doradus using the MeerKAT coverage. The results are shown in Figure 11.

We see that the transfer learned model reconstructs the image including the dynamic range well, despite not being trained on any data with a similar dynamic range. Since the number of measurements in the MeerKAT coverage is larger than that of the Gaussian coverage ($\sim 7.4\times$ more measurements), the reconstruction time is longer (194.6 ± 3.1 ms; $\sim 2.3\times$ longer) than for the simulated measurements and the training takes longer.




Figure 6. Reconstructions using the trained variants of the U-Net from simulated measurements of galaxy images using the true uv -coverage drawn from a Gaussian random distribution. Without any retraining, the single coverage and distribution of coverages strategies, struggle to reproduce the reconstruction quality achieved by fully retraining the models (the true coverage strategy), as indicated by the reconstruction SNR values. In particular the single coverage strategy performs poorly, sometimes even showing hallucinations of faint structure (e.g. the fourth example top centre). Using transfer learning improves the SNR of the reconstructions for both strategies, though performance is still lacking compared to the true coverage, oracle strategy.

5 CONCLUSIONS

Since the uv -coverage of every radio interferometric observation depends on factors such as the pointing direction of the telescope and the rotation of the Earth, it is essential to provide reconstruction techniques that generalise across varying visibility coverages. To this purpose we have proposed several training strategies for learned reconstruction methods that provide such robustness with minimal to no retraining of the reconstruction networks.

First of all, we conclude that methods that do not include the measurement operator in the reconstruction network, such as learned post-processing methods based on a U-Net model, are not robust to changes in the uv -coverage and are highly reliant on learning features that are part of the training data. While these methods are the most computationally efficient reconstruction approaches, their reconstructive performance is significantly worse than methods that do include the measurement operator, even when fully (re)trained or fine-tuned to the true uv -coverage, due to overfitting. Furthermore, by looking at the performance of this post-processing approach on more realistic radio interferometric observations, we have shown that these methods struggle to capture the high dynamic range of these images across the different training strategies. This indicates that

these methods are limited by the prior information in their training data, which is concerning in the context of real data where the true image is unknown.

On the other hand, we have shown that methods that include the measurement operator in the reconstruction network, such as the GU-Net model (Mars et al. 2023), maintain a high reconstruction quality across different uv -coverages with minimal to no need for retraining. In case of a Gaussian distribution of uv -coverages, training on a single different uv -coverages provides us with a network which suffers only a minor decrease in reconstruction quality on other uv -coverages of the same distribution with no retraining. With only a small amount of fine-tuning on the true uv -coverage, we come close to the reconstruction quality of a fully retrained network. This indicates that these methods are able to learn a more general representation of the prior information in the training data, irrespective of the uv -coverage.

Furthermore, we have shown that when applying the GU-Net reconstruction method to measurements outside of the training domain, such as a more realistic radio observation with a high dynamic range, we retain this reconstruction quality, indicative of the approach being




Figure 7. Reconstructions using the trained variants of the GU-Net from simulated measurements of galaxy images using the true uv -coverage drawn from a Gaussian random distribution. The strategies without retraining, the single coverage and distribution of coverages, generalise well to the unseen true coverage at only a minor loss in SNR compared to the oracle model trained on the true coverage. A small amount of fine-tuning using transfer learning brings the performance of these strategies closer to the true coverage, oracle model.




Figure 8. Reconstructions using the trained variants of the U-Net on simulated measurements, using the true uv -coverage drawn from a Gaussian random distribution, of an out-of-sample image of 30 Doradus, displayed on a linear scale (top) and a logarithmic scale (bottom). The strategies that are not fine-tuned on data from the true coverage perform significantly worse in reconstructing this out-of-sample image. Using transfer learning from the single coverage strategy brings the reconstructive SNR closer to fully retraining the model on the true coverage. All U-Net reconstructions struggle to accurately reconstruct the high dynamic range of this image as indicated by the logarithmically scaled images and the logSNR, with the single coverage strategy producing large artefacts.




Figure 9. Reconstructions using the trained variants of the GU-Net on simulated measurements, using the true uv -coverage drawn from a Gaussian random distribution, of an out-of-sample image of 30 Doradus, displayed on a linear scale (top) and a logarithmic scale (bottom). The GU-Net is able to reconstruct the high dynamic range of the image across the different strategies as indicated by both the high SNR and logSNR of the reconstructions. Training the model on a single coverage and then fine-tuning it using transfer learning provides reconstruction quality close to that of the model that is fully trained on simulated data from the true coverage.




Figure 10. The uv -coverage of an observation with the MeerKAT telescope.

significantly less reliant on the prior information in the training data and able to generalise across different types of data.

Finally, we have demonstrated that this unrolled iterative reconstruction method based on the GU-Net can also be used to reconstruct images from more realistic uv -coverage distributions, such as the MeerKAT uv -coverage, with high quality. While we show reconstruction models fine-tuned from Gaussian uv -coverage to the realistic MeerKAT uv -coverage, we recommend that future work investigate networks trained on a distribution of uv -coverages that is more representative of the true uv -coverage distribution, as this might provide us with a better starting point for fine-tuning the networks to the true uv -coverage.

In order to train these models on distributions of realistic uv -coverages, some technical challenges need to be addressed regarding the methods of varying the coverage during training causing a large overhead in terms of memory and/or compilation time. During the training of our networks on random Gaussian coverages, we have addressed this by using a random subset of a larger coverage, which

allows us to train on a distribution of coverages without the need to change the coverage during training, thus requiring only a minor increase in memory necessary. However, this would be highly computationally demanding in the case of realistic uv -coverages. In order to address this, we propose to investigate the use of approximate measurement operators in the reconstruction network by using a convolution with the PSF to approximate the measurement operator as a way to reduce the computational overhead of changing the coverage during training, as has been investigated for proximal optimisation methods (Bester et al. 2021). This would allow us to train on a distribution of realistic uv -coverages without the need to change the coverage during training, which would be essential in order to train these networks on realistic uv -coverages. However, the difference in performance between using the full measurement operator and the approximate operator in learned methods should be investigated further to assess how much information is lost.

Another interesting avenue for future work would be to investigate not only the fine tuning on the true uv -coverage, but also the performance of transfer learning on real data. In Mars et al. (2023) we have shown that transfer learning can be used to improve the performance of our reconstruction methods trained on a general image data set, by using a small amount of domain specific data. This approach could bridge the gap between a larger data set of simulated data and a small data set of real data.

Finally, a more extensive investigation should be done with regards to the performance of these methods on even more realistic radio observations, including beam errors, calibration errors and other errors that might be present in the actual data, as well as including bright point sources in the data sets which are not present in the current training data set.

In conclusion, we have presented a comprehensive study of the robustness of learned reconstruction methods to changes in the uv -coverage, and have shown that methods that include the measurement operator in the reconstruction network are able to generalise across different uv -coverages with minimal to no need for retraining, while maintaining a high reconstruction quality. We have shown that these methods are able to learn a generalised representation of the prior information in the training data, irrespective of the uv -coverage, and are able to generalise across different types of data, which is essential when looking to apply these methods to real data.




Figure 11. Reconstructions using a fine-tuned variant of the GU-Net on simulated measurements on an out-of-sample image of 30 Doradus using the MeerKAT uv -coverage as displayed in Figure 10. The ground truth, dirty image and GU-Net reconstruction are displayed on a linear scale (top) and a logarithmic scale (bottom). Using this fine-tuned model we achieve a high quality reconstruction that also reproduces the high dynamic range of the radio observation, using only minimal amount of re-training.

ACKNOWLEDGEMENTS

We thank Tariq Blecher for interesting discussions. Matthijs Mars is supported by the UCL Centre for Doctoral Training in Data Intensive Science (STFC Training grant ST/P006736/1). This work is also partially supported by EPSRC (grant number EP/W007673/1). The authors acknowledge the use of the UCL Myriad High Performance Computing Facility (Myriad@UCL), and associated support services, in the completion of this work. This work also used computing equipment funded by the Research Capital Investment Fund (RCIF) provided by UKRI, and partially funded by the UCL Cosmoparticle Initiative.

DATA AVAILABILITY

The code used to generate the results in this paper is available at <https://github.com/astro-informatics/LeIA>. The data used in this paper is available upon reasonable request to the corresponding author.

REFERENCES

- Adler J., Öktem O., 2018, *IEEE Transactions on Medical Imaging*, 37, 1322
- Allam Jr T., 2016, MSc. Thesis, University College London
- Beck A., Teboulle M., 2009, *SIAM Journal on Imaging Sciences*, 2, 183
- Bester H. L., Repetti A., Perkins S., Smirnov O. M., Kenyon J. S., 2021, A Practical Preconditioner for Wide-Field Continuum Imaging of Radio Interferometric Data (arxiv:2101.08072), doi:10.48550/arXiv.2101.08072
- Bhatnagar S., Cornwell T. J., 2004, *Astronomy & Astrophysics*, 426, 747
- Bhatnagar S., Cornwell T. J., Golap K., Uson J. M., 2008, *Astronomy & Astrophysics*, 487, 419
- Boink Y., Brune C., Manohar S., 2019, in Oraevsky A. A., Wang L. V., eds, *Photons Plus Ultrasound: Imaging and Sensing 2019*. SPIE, San Francisco, United States, p. 47, doi:10.1117/12.2507446
- Boink Y. E., Manohar S., Brune C., 2020, *IEEE Transactions on Medical Imaging*, 39, 129
- Boyd S., 2010, *Foundations and Trends® in Machine Learning*, 3, 1
- Briggs D. S., 1995, PhD thesis
- Cai X., Pereyra M., McEwen J. D., 2018a, *Monthly Notices of the Royal Astronomical Society*, 480, 4154
- Cai X., Pereyra M., McEwen J. D., 2018b, *Monthly Notices of the Royal Astronomical Society*, 480, 4170
- Carrillo R. E., McEwen J. D., Wiaux Y., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 1223
- Chambolle A., Ehrhardt M. J., Richtárik P., Schönlieb C.-B., 2018, *SIAM*

- Journal on Optimization*, 28, 2783
- Chen H., Zhang Y., Kalra M. K., Lin F., Chen Y., Liao P., Zhou J., Wang G., 2017, *IEEE Transactions on Medical Imaging*, 36, 2524
- Clark B. G., 1980, *Astronomy and Astrophysics*, 89, 377
- Connor L., Bouman K. L., Ravi V., Hallinan G., 2022, *Monthly Notices of the Royal Astronomical Society*, 514, 2614
- Cornwell T. J., 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 793
- Cornwell T. J., Golap K., Bhatnagar S., 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 647
- Dabbech A., Terris M., Jackson A., Ramatsoku M., Smirnov O. M., Wiaux Y., 2022, *The Astrophysical Journal Letters*, 939, L4
- Daubechies I., 1992, Ten Lectures on Wavelets. CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics, doi:10.1137/1.9781611970104
- Dewdney P., Hall P., Schilizzi R., Lazio T., 2009, *Proceedings of the IEEE*, 97, 1482
- Duijndam A., Schonewille M., 1997, in SEG Technical Program Expanded Abstracts 1997. Society of Exploration Geophysicists, pp 1135–1138, doi:10.1190/1.1885591
- Fessler J., Sutton B., 2003, *IEEE Transactions on Signal Processing*, 51, 560
- Gheller C., Vazza F., 2021, *Monthly Notices of the Royal Astronomical Society*, 509, 990
- Gregor K., LeCun Y., 2010, in Proceedings of the 27th International Conference on International Conference on Machine Learning. ICML'10. Omnipress, Madison, WI, USA, pp 399–406, doi:10.5555/3104322.3104374
- Hauptmann A., Adler J., Arridge S., Oktom O., 2020, *IEEE Transactions on Computational Imaging*, 6, 843
- Högbom J. A., 1974, *Astronomy and Astrophysics*, 15, 55
- Hotan A. W., et al., 2021, *Publications of the Astronomical Society of Australia*, 38, e009
- Jackson J., Meyer C., Nishimura D., Macovski A., 1991, *IEEE Transactions on Medical Imaging*, 10, 473
- Jin K. H., McCann M. T., Froustey E., Unser M., 2017, *IEEE Transactions on Image Processing*, 26, 4509
- Kingma D. P., Ba J., 2017, Adam: A Method for Stochastic Optimization (arxiv:1412.6980), doi:10.48550/arXiv.1412.6980
- Liaudat T. I., Mars M., Price M. A., Pereyra M., Betcke M. M., McEwen J. D., 2023, Scalable Bayesian Uncertainty Quantification with Data-Driven Priors for Radio Interferometric Imaging (arxiv:2312.00125), doi:10.48550/arXiv.2312.00125
- Maier A. K., et al., 2019, *Nature Machine Intelligence*, 1, 373
- Mars M., Betcke M. M., McEwen J. D., 2023, *RAS Techniques and Instruments*, 2, 760
- Nelson D., et al., 2019a, *Computational Astrophysics and Cosmology*, 6, 2
- Nelson D., et al., 2019b, *Monthly Notices of the Royal Astronomical Society*, 490, 3234
- Offringa A. R., et al., 2014, *Monthly Notices of the Royal Astronomical Society*, 444, 606
- Pratley L., Johnston-Hollitt M., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 3483
- Pratley L., McEwen J. D., d’Avezac M., Carrillo R. E., Onose A., Wiaux Y., 2018, *Monthly Notices of the Royal Astronomical Society*, 473, 1038
- Pratley L., McEwen J. D., d’Avezac M., Cai X., Perez-Suarez D., Christidi I., Guichard R., 2019a, Distributed and Parallel Sparse Convex Optimization for Radio Interferometry with PURIFY (arxiv:1903.04502), doi:10.48550/arXiv.1903.04502
- Pratley L., Johnston-Hollitt M., McEwen J. D., 2019b, *The Astrophysical Journal*, 874, 174
- Ronneberger O., Fischer P., Brox T., 2015, in Navab N., Hornegger J., Wells W. M., Frangi A. F., eds., Vol. 9351, Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. Springer International Publishing, Cham, pp 234–241, doi:10.1007/978-3-319-24574-4_28
- Ryu E. K., Liu J., Wang S., Chen X., Wang Z., Yin W., 2019, Plug-and-Play Methods Provably Converge with Properly Trained Denoisers (arxiv:1905.05406), doi:10.48550/arXiv.1905.05406
- Sault R. J., Staveley-Smith L., Brouw W. N., 1996, *Astronomy and Astrophysics Supplement Series*, 120, 375
- Schwab F. R., 1984, *The Astronomical Journal*, 89, 1076
- Steer D. G., Dewdney P. E., Ito M. R., 1984, *Astronomy and Astrophysics*, 137, 159
- Stewart I. M., Fenech D. M., Muxlow T. W. B., 2011, *Astronomy & Astrophysics*, 535, A81
- Tasse C., van der Tol S., van Zwieten J., van Diepen G., Bhatnagar S., 2013, *Astronomy and Astrophysics*, 553, A105
- Taylor G. B., Carilli C. L., Perley R. A., Observatory N. R. A., eds, 1999, Synthesis Imaging in Radio Astronomy. 2: Sixth NRAO/NMIMT Synthesis Imaging Summer School, 17 - 23 June 1998 / Ed. by G. B. Taylor, C. L. Carilli, and R. A. Perley. No. Vol. 180 in Conference Series / Astronomical Society of the Pacific, Astronomical Society of the Pacific, San Francisco, Calif
- Terris M., Abdulaziz A., Dabbech A., Jiang M., Repetti A., Pesquet J.-C., Wiaux Y., 2019, Signal Processing with Adaptive Sparse Structured Representations (SPARS) workshop, p. 3
- Terris M., Dabbech A., Tang C., Wiaux Y., 2022, *Monthly Notices of the Royal Astronomical Society*, 518, 604
- Terris M., Tang C., Jackson A., Wiaux Y., 2023, Plug-and-Play Imaging with Model Uncertainty Quantification in Radio Astronomy (arxiv:2312.07137), doi:10.48550/arXiv.2312.07137
- Thompson A. R., Moran J. M., Swenson G. W., 2017, Interferometry and Synthesis in Radio Astronomy. Astronomy and Astrophysics Library, Springer International Publishing, Cham, doi:10.1007/978-3-319-44431-4
- Tingay S. J., et al., 2013, *Publications of the Astronomical Society of Australia*, 30, e007
- Trent D. B., 2020, MSc. Thesis, University College London, London
- Van Haarlem M. P., et al., 2013, *Astronomy & Astrophysics*, 556, A2
- Venkatakrishnan S. V., Bouman C. A., Wohlberg B., 2013, in 2013 IEEE Global Conference on Signal and Information Processing. IEEE, Austin, TX, USA, pp 945–948, doi:10.1109/GlobalSIP.2013.6737048
- Wilber A. G., Dabbech A., Terris M., Jackson A., Wiaux Y., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 5576
- Yi X., Babyn P., 2018, *Journal of Digital Imaging*, 31, 655

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.