

PolarRec: Improving Radio Interferometric Data Reconstruction Using Polar Coordinates

Ruoqi Wang¹, Zhuoyang Chen¹, Jiayi Zhu¹, Qiong Luo^{*,1,2}, Feng Wang³

¹ The Hong Kong University of Science and Technology (Guangzhou),

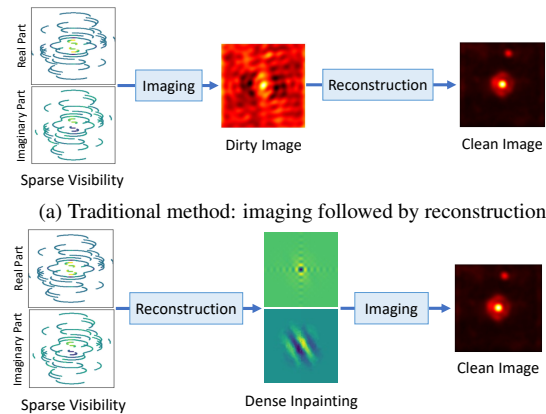
²The Hong Kong University of Science and Technology, ³ Guangzhou University
 {rwang280, zchen190, jzhu718}@connect.hkust-gz.edu.cn, luo@ust.hk, fengwang@gzhu.edu.cn

Abstract

In radio astronomy, visibility data, which are measurements of wave signals from radio telescopes, are transformed into images for observation of distant celestial objects. However, these resultant images usually contain both real sources and artifacts, due to signal sparsity and other factors. One way to obtain cleaner images is to reconstruct samples into dense forms before imaging. Unfortunately, existing reconstruction methods often miss some components of visibility in frequency domain, so blurred object edges and persistent artifacts remain in the images. Furthermore, the computation overhead is high on irregular visibility samples due to the data skew. To address these problems, we propose PolarRec, a transformer-encoder-conditioned reconstruction pipeline with visibility samples converted into the polar coordinate system. This coordinate system matches the way in which radio telescopes observe a celestial area as the Earth rotates. As a result, visibility samples distribute in the polar system more uniformly than in the Cartesian space. Therefore, we propose to use radial distance in the loss function, to help reconstruct complete visibility effectively. Also, we group visibility samples by their polar angles and propose a group-based encoding scheme to improve the efficiency. Our experiments demonstrate that PolarRec markedly improves imaging results by faithfully reconstructing all frequency components in the visibility domain while significantly reducing the computation cost in visibility data encoding. The code is available at <https://github.com/RapidsAtHKUST/PolarRec>.

1. Introduction

In radio astronomy, *visibility* refers to radio signal data from celestial objects, obtained by radio telescopes. These data are represented as complex values in the *uv-plane*, a geometric plane defined for interferometric observations. Vis-



(a) Traditional method: imaging followed by reconstruction

(b) Recent method: reconstruction followed by imaging.

Figure 1. Two visibility data processing flows.

ibility data are subsequently converted into images through *imaging* for further analysis. However, these images, known as *dirty images*, are often dominated by artifacts [23]. This phenomenon is due to limitations in telescope configurations, under which not the entire *uv-plane* is sampled. Therefore, visibility data normally have to be reconstructed before being utilized in scientific analysis. In this paper, we propose a visibility reconstruction method, aiming to reconstruct the real sky by recovering all visibility components in the *uv-plane* effectively and efficiently.

Traditional methods first transfer the sparse visibility data into dirty images through the imaging process and then reconstruct the dirty images to *clean images* [1, 3, 10, 14, 26, 30]. The process is shown in Figure 1 (a). In contrast, some recent deep-learning-based studies [23, 31] have proposed to first do inpainting on the visibility data to reconstruct the sparse samples to dense coverage and then perform imaging to obtain the clean image, as shown in Figure 1 (b), showing better performance. In order to produce a high-fidelity sharp image, it is crucial to densely sample the full visibility domain [31]. In this paper, we adopt the reconstruction-and-imaging processing flow.

Existing methods for reconstructing visibility data face

*Corresponding Author.




Figure 2. Visibility data samples with different coordinate systems.

challenges in both effectiveness and efficiency. Specifically, the first challenge is to effectively capture visibility components within the uv -plane. For example, Radionets [23] and U-Nets [22] use a convolutional neural network that is based on pixels and grids, resulting in discontinuities in reconstructed visibility maps. In comparison, Wu et al. [31] use neural fields to address the continuity problem, but their method can only restore the low-frequency part of visibility (located near the center of the uv -plane), largely missing the high-frequency portion (found far from the uv -plane center). Such discontinuity and incompleteness in the visibility domain cause blurred edges of observed objects, disappearance of faint astronomical sources, and persistence of artifacts, in resultant images.

The second major issue is the inefficiency of current strategies. For instance, Radionets [23] encode visibility data together with an excessive amount of zero-inpainting, causing a large amount of unnecessary computation. Wu et al. [31] embed each visibility sample point as a token in their Transformer encoder [29] to attend to each other, which is of quadratic computation cost in the number of sample points [12], too high for real applications. These limitations force a compromise between computational cost and quality of reconstruction in current methods, severely limiting the application of deep learning techniques in radio astronomy imaging and reconstruction.

To address these challenges, we explore the observational nature of radio telescopes and then propose PolarRec, which leverages the polar coordinates of sample points to enhance the reconstruction quality and reduce the cost. Our key observation is that visibility samples are obtained by telescopes as the Earth rotates, and that the high- and low-frequency signals are distributed on the uv -plane in accordance with their distances from the center. As a result, the visibility points distribute more evenly in the polar coordinate system than in Cartesian coordinates [19] (see Figure 2). Moreover, the polar-based system has shown promising properties in vision tasks including instance segmentation and object detection [7, 19, 32, 33]. Therefore, we propose to adopt the polar coordinate system, using the radial coordinate to associate with the frequency information in the

loss function and the angular coordinate to group sample points.

Specifically, we first design a loss function for visibility data with a weighting scheme based on the radial coordinates. Our weighted loss function, Radial Visibility Loss, associates visibility data with radial coordinates in the frequency domain, enabling effective reconstruction of visibility data, on both high- and low-frequency components. In comparison, existing work recovers mainly low-frequency components. Consequently, our method produces sharper, more detailed imaging results. Furthermore, by introducing group-level encoding of visibility points according to angular coordinates, our method is more efficient than visibility-point-level encoding. This group encoding improves computation efficiency, making the use of Transformer encoders for visibility encoding more practical and scalable.

In summary, our main contributions are as follows:

- We propose PolarRec to seamlessly integrate visibility data with a Transformer-encoded reconstruction pipeline by using polar coordinates, addressing two fundamental challenges in visibility reconstruction.
- We introduce the Radial Visibility Loss (RVL), which incorporates the radial coordinates of visibility data, adeptly capturing both low- and high-frequency visibility components.
- As the first to study the granularity of visibility encoding, we innovatively utilize angular coordinates to group sampled points for Transformer encoding, markedly enhancing computational efficiency.

We have experimentally evaluated our proposed PolarRec on public datasets of four galaxy morphologies, including an overall comparison with other state-of-the-art methods, ablation studies, and tests on group sizes and grouping techniques. The experimental results demonstrate the effectiveness and efficiency of our method.

2. Background and Related Work

2.1. Very Long Baseline Interferometry (VLBI)

In radio astronomy, using radio interference signals to image distant astronomical sources requires telescopes of very large aperture [4], because the angular resolution of a telescope is inversely proportional to the diameter. A major observation technique is the Very Long Baseline Interferometry (VLBI), which uses multiple radio telescopes spreading over the globe to form a virtual Earth-sized telescope. The radio waves from astronomical sources are recorded separately at individual telescopes. Then, these signals are cross-correlated for all pairs of antennas at a central location, generating *visibility* data. A VLBI observation is typically performed for hours to measure as many points in the uv -plane as possible. However, the measurement results remain sparse due to the limited number of antennas [4, 28].




Figure 3. Effects of band limiting.

Consequently, sparse-to-dense reconstruction on visibility data is necessary to improve the imaging quality for the following analysis.

2.2. Interferometric Imaging

Visibility data, represented as complex values, is the result of a Fourier transform of the sky’s brightness distribution [18]. *Imaging* converts visibility data into images, which can be analyzed to provide insights about the observed celestial bodies [28]. In the imaging process, an inverse Fourier transform maps the (u, v) coordinates from the Fourier domain to (l, m) coordinates in the image domain [31]. The transformation can be described as:

$$I(l, m) = \int_u \int_v e^{2\pi i(ul+vm)} V(u, v) du dv. \quad (1)$$

In this equation, $V(u, v)$ is the visibility data in Fourier space, and $I(l, m)$ represents the intensity distribution in the image domain.

2.3. Radio Interferometric Data Reconstruction

Radio interferometric data reconstruction is vital for converting sparse visibility data into clear, artifact-free images. However, previous methods in this domain [1, 3, 10, 14, 23, 26, 30, 31] have been limited by either their inability to capture continuous and complete frequency components or computational inefficiencies. To the best of our knowledge, no previous studies have managed to reconstruct all components in visibility data or increase the granularity of visibility encoding for efficiency.

3. Our Method

In this section, we first investigate the relation between visibility components and imaging results. After that, we present PolarRec, which adopts the polar coordinate system in visibility reconstruction. In PolarRec, we design Radial Visibility Loss to incorporate a weighting scheme based on the radial coordinate within the uv -plane. Moreover, we group the visibility points by their angular coordinates and then extract grouping tokens for the subsequent

Transformer encoding. An overview of our method is presented in Figure 4.

3.1. Imaging Results of Visibility

In radio interferometry, visibility data plays a vital role in generating high-fidelity images of celestial objects. The imaging process is the inverse Fourier transformation of visibility data to construct the brightness distribution of the observed sky. Consequently, missing or incomplete visibility components can significantly degrade the imaging result, leading to artifacts and loss of crucial information about the object’s structure. Therefore, we first investigate the impact of missing visibility components on the final imaging output.

We explore the impact of visibility components in various frequency regions by applying standard band-limiting operations [17] and analyzing their effects on the imaging results. As shown in Figure 3, the loss of high frequency data due to a low-pass filter results in blur and artifacts and causes the disappearance of weak sources. Meanwhile, a high-pass filter eliminates low-frequency data and reduces the overall quality, but it maintains clear object edges as well as dim or small sources. Additionally, band-and-stop filters lead to blurred images and artifacts.

In summary, different frequency ranges cause distinct imaging effects. This observation suggests that recovering all missing visibility components could enhance the quality of imaging results. More specifically, if a resultant image has blurred edges or misses dim light sources, it is probably due to poor reconstruction of the high-frequency components of the visibility data.

3.2. Polar Coordinates

Our utilization of the polar coordinates is to represent the visibility more uniformly. It is because visibility sampling is based on Earth’s rotation, and high- and low-frequency visibility components are distributed on the uv -plane according to their distances from the origin of the plane.

In the uv -plane, we convert (u, v) coordinates to polar coordinates, denoted as $(r(u, v), \theta(u, v))$:

$$r(u, v) = \sqrt{u^2 + v^2} \quad (2)$$

$$\theta(u, v) = \arctan2(v, u) \quad (3)$$

Where $r(u, v)$ represents the radial distance from the origin, and $\theta(u, v)$ represents the angle of the vector from the positive u -axis.

3.3. Radial Visibility Loss

We propose Radial Visibility Loss (RVL), incorporating the radial coordinates of visibility data to capture both low- and high-frequency visibility.

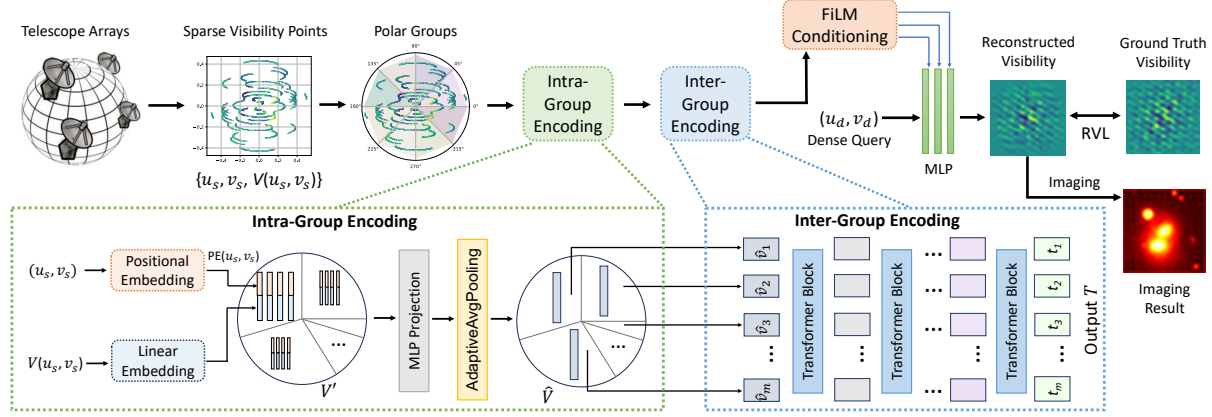


Figure 4. An overview of our method. Sparse visibility data $V(u_s, v_s)$ are grouped by the angular coordinate and passed through two encoding layers: intra-group encoding to generate group tokens and inter-group encoding by a Transformer encoder. The encoded output T then conditions the predicted visibility generation in the neural field. The final output is compared with the ground truth to compute the Radial Visibility Loss.

To compute RVL, we first compute the weight matrix $w_1(u, v)$ to down-weight easy visibility components (components whose predicted values are close to the ground truth) based on Focal Frequency Loss (FFL) [17]:

$$w_1(u, v) = |V_r(u, v) - V_p(u, v)|^\alpha \quad (4)$$

where $V_r(u, v) = A_r + iB_r$ and $V_p(u, v) = A_p + iB_p$ are the ground truth and predicted visibilities in complex form respectively, and α is a scaling factor.

We then introduce an additional weight $w_2(u, v)$, computed from $r(u, v)$, to make our model pay more attention to high-frequency components of the visibility during reconstruction. The weight term $w_2(u, v)$ is calculated as follows:

$$w_2(u, v) = \left(\frac{r(u, v)}{\max(r(u, v))} + 1 \right)^\beta \quad (5)$$

Where $\frac{r(u, v)}{\max(r(u, v))}$ normalizes the radial coordinate, ensuring the weight is more significant for points farther from the center, corresponding to higher frequencies in the visibility data. Adding 1 prevents any weights from becoming zero, maintaining the influence of all visibility components during the learning process. β is a scaling factor.

The final weight $w(u, v)$ is computed as:

$$w(u, v) = \left(\frac{r(u, v)}{\max(r(u, v))} + 1 \right)^\beta |V_r(u, v) - V_p(u, v)|^\alpha \quad (6)$$

A thorough visibility distance measure must consider both amplitude and phase as they hold distinct aspects of imaging information [17]. The complex visibility value can be represented as a vector within a two-dimensional plane. Specifically, we map ground truth visibility $V_r(u, v) =$

$A_r + iB_r$ and the corresponding predicted value $V_p(u, v) = A_p + iB_p$ to vectors \vec{p}_r and \vec{p}_p (as shown in Figure 5). According to the definition of 2D discrete Fourier transform, the vector magnitude relates to the amplitudes $|\vec{p}_r|$ and $|\vec{p}_p|$, and the angles θ_r and θ_p reflect the phase. Consequently, following FFL [17], we define the single-point visibility distance considering both magnitude and direction as:

$$d(\vec{p}_r, \vec{p}_p) = \|\vec{p}_r - \vec{p}_p\|_2^2 = |V_r(u, v) - V_p(u, v)|^2. \quad (7)$$

The final visibility distance is computed as the average of these individual distances over all visibility points:

$$d(V_r, V_p) = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} |V_r(u, v) - V_p(u, v)|^2. \quad (8)$$

The final expression for the Radial Visibility Loss is derived by executing the Hadamard product between the weight matrix $w(u, v) = w_1(u, v)w_2(u, v)$ and the visibility distance matrix $d(V_r, V_p)$, as follows:

$$\text{RVL} = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} w(u, v) |V_r(u, v) - V_p(u, v)|^2 \quad (9)$$




Figure 5. Example of visibility distance.

3.4. Encoding by Angular-Coordinate Groups

Using Transformers for encoding visibility samples causes quadratic computation cost in the number of sample points [12]. Such high costs constrain practical applications, especially as the number of visibility points in each image expands with more telescope arrays or extended observation duration. To solve the problem, we propose a group-based encoding scheme to improve the efficiency.

As shown on the left in Figure 4, the sparsely sampled visibility data is in the form of $\{u_s, v_s, V(u_s, v_s)\}$, where (u_s, v_s) are the coordinates at which a measurement is sampled and $V(u_s, v_s)$ is the complex value of the sample point. We embed the real and imaginary components in $V(u_s, v_s)$ separately.

First, we divide all sample points into m groups according to their angular coordinate $\theta(u_s, v_s)$. To integrate each visibility value and its corresponding position, we encode each sample individually using positional embedding ($\text{PE}(u_s, v_s)$ in Figure 4). Specifically, we encode the positional information of a sample point using Random Fourier Embedding [27]. After that, the positional embedding $\text{PE}(u_s, v_s)$ and the embedding of complex value of visibility $V(u_s, v_s)$ are concatenated to form visibility tokens V' . Denote the sparse visibility tokens as $V' = [v'_1; v'_2; v'_3; \dots; v'_n]$, where $V' \in \mathbb{R}^{n \times d}$, n is the number of sample measurement points and d is the number of dimensions of visibility tokens.

Then, we apply a Multi-Layer Perceptron (MLP) mapping layer and averaging aggregating for intra-group encoding. The encoding result is $\hat{V} = [\hat{v}_1; \hat{v}_2; \hat{v}_3; \dots; \hat{v}_m]$, where $\hat{V} \in \mathbb{R}^{m \times d}$. For each group i , we compute the group token \hat{v}_i as follows:

$$\hat{v}_i = \text{Avg}(\text{MLP}(v_j)), i = 1 \text{ to } m, v_j \in \text{group } i. \quad (10)$$

In the implementation of intra-group encoding, we sort the tokens V' of sample points according to their angular coordinates. After mapping these tokens through MLP, we use adaptive average pooling to compute the group tokens \hat{V} . For data collected with the same telescope configuration, the process of sorting based on angular coordinates needs to be performed only once, as the locations of sample points remain unchanged. The intra-group encoding process is as shown in Algorithm 1.

Finally, the group tokens \hat{V} go through inter-group encoding by a transformer encoder. We base our encoder design on Transformer structures similar to prior work [12, 29, 31]. The input group tokens are then transformed into latent tokens by multi-headed self-attention layers.

3.5. Neural Field Conditioning

Our method follows the conditional neural field pipeline proposed by Wu et al. [31]. Given the sparsely sampled

Algorithm 1 Intra-Group Encoding

```

1: for each  $u_{si}, v_{si}, V(u_{si}, v_{si})$  in set of sampled points do
2:    $\text{pe}_i \leftarrow \text{PE}(u_{si}, v_{si})$ 
3:    $\text{le}_i \leftarrow \text{LinearEmbedding}(V(u_{si}, v_{si}))$ 
4:    $v'_i \leftarrow \text{concat}(\text{pe}_i, \text{le}_i)$ 
5:    $\theta_i \leftarrow \text{arctan2}(u_{si}, v_{si})$ 
6: end for
7:  $V' \leftarrow [v'_1; v'_2; \dots; v'_n], \Theta \leftarrow [\theta_1; \theta_2; \dots; \theta_n]$ 
8: if sorted_indices not precomputed then
9:   sorted_indices  $\leftarrow \text{argsort}(V', \Theta)$ 
10: end if
11:  $V'' \leftarrow V'[\text{sorted\_indices}]$ 
12:  $V^{\text{MLP}} \leftarrow \text{MLP}(V'')$ 
13:  $\hat{V} \leftarrow \text{AdaptiveAvgPooling}(V^{\text{MLP}}, m)$ 
14: return  $\hat{V}$ 

```

visibility $V(u_s, v_s)$, our objective is to determine a neural field Φ , fulfilling a constraint set by the function F :

$$F(\Phi(u_s, v_s), V(u_s, v_s)) = 0. \quad (11)$$

We approximate this implicit function $\Phi(u, v)$ with an MLP of l layers parameterized by weights Θ_m .

We use the output tokens of inter-group encoding $T = [t_1; t_2; t_3; \dots; t_l]$ to extend the neural field with a learning-based prior, with each token corresponding to an MLP layer. Using the FiLM conditioning [21], the output tokens modulate the i th layer’s activation \mathbf{x}_i by:

$$\text{FiLM}(\mathbf{x}_i) = \gamma(t_i) \odot \mathbf{x}_i + \beta(t_i), i \in 1 \text{ to } l, \quad (12)$$

where γ and β are simple affine layers with non-linearities and \odot signifies a Hadamard product [16].

The MLP parameters Θ_m and the encoder parameters Θ_e are jointly optimized during training:

$$\min_{\Theta_m, \Theta_e} \text{RVL}(\Phi(u_d, v_d; \{T\}; \Theta_m), V_{\text{gt}}(u_d, v_d)), \quad (13)$$

with $\{T\} = \Psi(\{u_s, v_s, V(u_s, v_s)\}; \Theta_e)$,

where (u_d, v_d) are the dense coordinates in visibility plane and $V_{\text{gt}}(u_d, v_d)$ is the ground truth of visibility inpainting.

4. Experiments

In this section, we conduct a comprehensive evaluation of our method in comparison with several classic and recent state-of-the-art methods to demonstrate the overall improvement achieved by our approach. We also design experiments to explore the effects of different grouping methods and group sizes on the reconstruction results. In addition, we conduct ablation experiments to study the effects of individual weighting techniques in RVL.

4.1. Experimental Setup

Platform. We conduct all experiments on a server with two AMD EPYC 7302 CPUs, 128GB main memory, and eight

Nvidia RTX 3090 GPUs each with 24GB device memory. The server is equipped with an NVME 2TB SSD and four 1TB SATA hard disks. The operating system is Ubuntu 20.04. Our model is implemented in PyTorch 1.8.1 [20].

Datasets. In the experiments, we evaluate our method on different public datasets of distinct galaxy morphologies [8, 13]: Merging Galaxies (MG), In-between Round Smooth Galaxies (IRSG), Unbarred Tight Spiral Galaxies (UTSG), and Edge-on Galaxies with Bulge (EGB). These data are derived from sources under the DESI Legacy Imaging Surveys [11], integrating contributions from the Beijing-Arizona Sky Survey (BASS) [35], the DECam Legacy Survey (DECaLS) [2], and the Mayall z-band Legacy Survey [24]. The visibility data are generated from these images using the eht-imaging toolkit [5, 6], denoted by $\{u_s, v_s, V(u_s, v_s)\}$. The parameters for observation are adjusted to mirror an 8-telescope Event Horizon Telescope (EHT) setup [31], with the EHT being one of the most prominent arrays leveraging VLBI techniques. All visibility data generated from our image datasets are in 128X128 grids. We sample the Fourier transform results of our image datasets to use as the ground truth. In each image dataset, we randomly select 80% of the images for training, and the other 20% for testing.

Methods under Comparison. We compare our method with three other methods for radio interferometry reconstruction, including the classic method CLEAN [14], which is for dirty image reconstruction, and two latest deep learning-based approaches for visibility data reconstruction – Radionets [23] and Neural Interferometry [31]. We use the original code of these methods and follow the parameter setting in the original code for the best performance. All these methods are implemented on PyTorch. In addition, we test the magnetic resonance imaging (MRI) reconstruction U-Net [22, 34] to reconstruct visibility data as supplementary baselines. Compared with CLEAN [14] which assumes point sources, we use learning-based methods to learn priors from data for more accurate reconstruction. Different from leaning-based methods [22, 31, 34], PolarRec utilizes polar coordinates for visibility data. We use the polar groups as encoding granularity to improve the efficiency whereas others use pixel grids [22, 34] or points [31]. We use radial coordinate based loss, RVL, which can capture both low- and high-frequency visibility. In contrast, others use simple MSE [22, 31, 34].

Evaluation Metrics. To measure differences in frequency data, we use the Log Frequency Distance (LFD) [17], which is defined as follows:

$$\text{LFD} = \log \left[\frac{1}{MN} \left(\sum_{u=0}^{M-1} \sum_{v=0}^{N-1} |V_r(u, v) - V_p(u, v)|^2 \right) + 1 \right] \quad (14)$$

where $V_r(u, v)$ and $V_p(u, v)$ represent the real and predicted visibility respectively. A lower LFD is better.

To evaluate the quality of images produced from the




Figure 6. Visual comparison of visibility and image. The reconstructed images, visibility data and their corresponding error maps are provided. The small panels to the left of each image are the corresponding visibility data, with the real part on top and the imaginary part at the bottom.

reconstructed visibility, we employ two common metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). PSNR quantifies the overall image quality of the resultant images, and SSIM quantifies the perceptual similarity to the ground truth images. We compute these two metrics using the scikit-image package [25], which follows the formulas presented by Hore et al. [15]. A higher PSNR and SSIM is better.

To evaluate the efficiency, we use inference time and Floating Point Operations (FLOPs). The inference time in the experiments is tested on a single Nvidia RTX 3090. We compare time performance between our method and Neural Interferometry only, because only these two are transformer-encoder based.

4.2. Overall Comparison

We calculate the LFD, PSNR, and SSIM values for all test images reconstructed using our method and other methods, presenting both mean values and standard deviations. These results are listed in Table 1. The results show that PolarRec consistently outperforms the other methods in all three measures and four datasets, underscoring the effectiveness of our reconstruction method.

We also illustrate some representative reconstructed visibility and the corresponding images on the four datasets in Figure 6, including all deep-learning-based methods under comparison as well as the dirty images and the ground truth images of the real sky. Comparing the dirty images in Figure 6 (a) and the ground truth in Figure 6 (f), we find there

Table 1. Overall performance comparison (mean and standard deviation) on different datasets. All comparison methods have significant difference with our method ($p < 0.01$).

Models	Img Vis	MG				IRSG				UTSG				EGB			
		LFD ↓	PSNR ↑	SSIM ↑		LFD ↓	PSNR ↑	SSIM ↑		LFD ↓	PSNR ↑	SSIM ↑		LFD ↓	PSNR ↑	SSIM ↑	
Dirty		N/A	10.453 (1.256)	0.680 (0.055)		N/A	10.875 (1.293)	0.711 (0.050)		N/A	10.655 (1.092)	0.690 (0.047)		N/A	10.158 (1.216)	0.674 (0.055)	
CLEAN [14]	✓	N/A	18.708 (2.336)	0.767 (0.031)		N/A	20.974 (2.192)	0.795 (0.025)		N/A	17.200 (2.465)	0.750 (0.035)		N/A	20.175 (2.411)	0.779 (0.033)	
U-Net [34]	✓	1.171 (0.198)	18.126 (1.906)	0.818 (0.021)		1.235 (0.230)	19.770 (1.925)	0.828 (0.024)		1.256 (0.227)	14.877 (1.701)	0.786 (0.032)		1.219 (0.234)	19.232 (2.078)	0.822 (0.024)	
Radionets [23]	✓	1.580 (0.202)	19.687 (1.897)	0.836 (0.022)		1.499 (0.206)	21.369 (1.887)	0.854 (0.022)		1.419 (0.231)	20.828 (2.003)	0.844 (0.022)		1.545 (0.220)	20.322 (1.809)	0.836 (0.023)	
Neu.Int [31]	✓	1.229 (0.282)	20.560 (2.287)	0.875 (0.028)		0.916 (0.302)	24.099 (2.760)	0.898 (0.027)		1.038 (0.313)	21.323 (2.303)	0.880 (0.027)		1.146 (0.343)	21.276 (2.675)	0.879 (0.031)	
PolarRec(Ours)	✓	0.924 (0.284)	24.206 (2.547)	0.889 (0.029)		0.666 (0.238)	26.540 (2.653)	0.911 (0.025)		0.707 (0.241)	25.430 (2.565)	0.905 (0.025)		0.707 (0.241)	25.880 (2.529)	0.906 (0.025)	

are many artifacts and distortion of object structure in dirty images because of the sparsity of the visibility.

As shown in Figure 6 (b), reconstructions of the visibility by MRI reconstruction U-Net [34] are dominated by the artifacts in resultant images. In Figure 6 (c), Radionets [23] are able to reconstruct more visibility content than the U-Net. However, the reconstruction is discontinuous. Although Radionets reduce artifacts in the imaging results, it cannot distinguish between separate sources that are in close proximity. Furthermore, many faint astronomical sources are missing in the reconstructed images. In contrast, as shown in Figure 6 (d), Neural Interferometry [31] (denoted Neu-Int) can continuously and realistically reconstruct the low-frequency components of the visibility, but misses much information from the high-frequency components, leading to a loss of details in the reconstruction.

The results of PolarRec (Figure 6 (e)) show that our method can effectively reconstruct more complete and continuous visibility data than others. The resultant imaging results not only eliminate artifacts but also restore the true structure of astronomical sources while preserving details and small faint sources.

4.3. Computational Cost Comparison

To demonstrate the effectiveness of the group-granularity encoding in PolarRec, we conduct a comparative analysis against the Neural Interferometry [31] method with point-granularity encoding due to both methods are built upon the Transformer architecture. We fix the number of visibility points as 1660 and the size of the reconstructed visibility map as 128×128 , and compare the computational efficiency of the encoders with batch size from 4 to 32. The comparison focuses on two main metrics: the number of Floating Point Operations (GFLOPs) and inference time (Latency (ms)).

Table 2 shows the results of this comparison. The latency of group granularity encoding in PolarRec is significantly lower than point granularity encoding in Neural Interferometry [31]. Moreover, as the batch size increases, the rise in latency for PolarRec is much smaller compared to Neural Interferometry [31]. These results indicate that

Table 2. Comparative Analysis of Computational Efficiency

Method	BatchSize	Latency (ms)	GFLOPs
Neural Interferometry	4	29.53 (3.33)	44.02
	8	59.02 (3.19)	88.04
	16	115.62 (7.30)	176.08
	32	232.04 (5.64)	352.16
PolarRec Group Size = 32	4	3.72 (1.84)	3.05
	8	3.89 (3.23)	6.10
	16	4.20 (2.00)	12.20
	32	5.37 (3.05)	24.40
PolarRec Group Size = 64	4	3.47 (0.10)	2.39
	8	3.90 (2.94)	4.78
	16	4.02 (2.11)	9.57
	32	5.11 (1.91)	19.14

PolarRec has greater advantages in handling large batches of data, which has important implications for the processing and analysis of massive astronomical observational data.

4.4. Effect of Group Size

In this experiment, we vary the group size in number of visibility points and investigate its impact on the reconstruction results. We also record the inference latency of the encoding process, and use Floating Point Operations (FLOPs) to measure the computation cost. LFD, PSNR and SSIM are used to assess the reconstruction quality.

As illustrated in Table 3, there is a sharp drop in both FLOPs and inference latency as the group size increases from 1 to 16. In contrast, the image quality in SSIM and PSNR is almost constant under most settings. Only the PSNR value of MG dataset decreases slightly. Between a group size of 32 to 128, there is a slight decrease in PSNR and SSIM, implying that increasing group size beyond 32

Table 3. Effects of group size. The experiment compares PSNR and SSIM metrics against GFLOPs and latency across different group sizes for four datasets.

Group Size	PSNR				SSIM				GFLOPs	Time (ms)
	MG	IRSG	UTSG	EGB	MG	IRSG	UTSG	EGB		
1	24.29	26.75	25.38	25.97	0.890	0.910	0.903	0.904	366.14	221.53
4	24.21	26.54	25.43	25.88	0.889	0.911	0.905	0.906	98.93	24.43
8	23.39	26.52	25.50	25.79	0.881	0.911	0.905	0.904	56.14	11.71
16	23.37	26.50	25.28	25.67	0.881	0.910	0.904	0.904	34.95	6.75
32	23.51	26.80	25.63	26.07	0.881	0.912	0.905	0.906	24.40	4.83
64	23.36	26.20	24.94	25.62	0.882	0.911	0.902	0.905	19.14	4.46
128	22.06	26.02	24.48	25.30	0.866	0.906	0.897	0.904	16.51	4.41

might compromise the output quality. Moreover, when the group size is set to 1, it is the same as encoding at the granularity of individual points. The results indicate that encoding with group granularity as input to the transformer encoder is significantly more efficient than encoding at the point granularity even when the group size is small.

4.5. Ablation Experiment

This ablation experiment aims to examine the significance of two weight matrices w_1, w_2 in the Radial Visibility Loss by omitting them one at a time. The results in Table 4 and Figure 7 show that the full Radial Visibility Loss is the best. Removing either component w_1 , or w_2 results in reduced performance across all metrics.

Table 4. Ablation study performance comparison (mean and standard deviation) on different datasets.

	MG			IRSG		
	LFD ↓	PSNR ↑	SSIM ↑	LFD ↓	PSNR ↑	SSIM ↑
w/o w_2	1.088 (0.309)	22.498 (2.578)	0.884 (0.030)	0.876 (0.305)	24.435 (2.577)	0.904 (0.026)
w/o w_1	1.075 (0.317)	22.706 (2.797)	0.882 (0.032)	0.818 (0.283)	25.213 (2.641)	0.907 (0.027)
RVL	0.924 (0.284)	24.206 (2.547)	0.889 (0.029)	0.666 (0.238)	26.540 (2.653)	0.911 (0.025)
	UTSG			EGB		
	LFD ↓	PSNR ↑	SSIM ↑	LFD ↓	PSNR ↑	SSIM ↑
w/o w_2	0.908 (0.299)	23.477 (2.574)	0.897 (0.027)	0.933 (0.310)	23.793 (2.507)	0.897 (0.026)
w/o w_1	0.854 (0.281)	24.393 (2.627)	0.901 (0.028)	0.864 (0.287)	24.569 (2.500)	0.901 (0.026)
RVL	0.707 (0.241)	25.430 (2.565)	0.905 (0.025)	0.707 (0.241)	25.880 (2.529)	0.906 (0.025)




Figure 7. Visual examples of ablation experiments.

4.6. Effect of Grouping Method

We also vary the grouping method and measure its performance impact on PolarRec. We implemented three grouping strategies: (1) grouping by clustering visibility points based on their positions, (2) grouping by the radial coordinate, and (3) grouping by the angular coordinate. Each of these strategies was tested under various group sizes on different datasets. The LFD and PSNR results are presented in Figure 8. Regardless of the group size, grouping the visibility points by angular coordinates always has the best performance.

4.7. Test on Real-World Observations

We test our model trained on the Galaxy10 synthetic dataset [13] with real observational data [9]. An example is shown

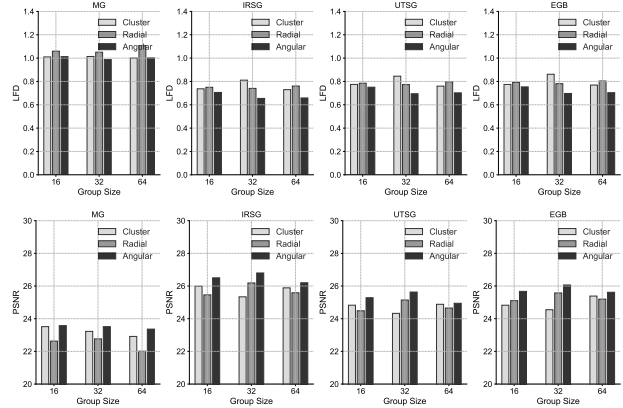


Figure 8. Comparison of grouping method over various group sizes and different datasets.

in Figure 9. Our reconstructed image is of much higher quality than the dirty image from the original sparse visibility data even though some artifacts remain. Our future work will focus on improving the generalizability of our visibility reconstruction methods.




Figure 9. An Example of Reconstruction of Real Data

5. Conclusion and Future Work

We have presented PolarRec, reconstructing interferometric visibility with sample points represented in the polar coordinate system. By adopting angular and radial coordinates of visibility points, our method addresses two critical challenges in visibility reconstruction: reconstruction quality and computational efficiency. Our results show that PolarRec markedly improves imaging outcomes while significantly reducing the computation cost, demonstrating a new, feasible, and extendable approach for cross-disciplinary progress in computer vision and radio astronomy.

While PolarRec is designed for radio astronomical data reconstruction, its underlying principles may find utility in other domains requiring sophisticated frequency component analysis and reconstruction, such as magnetic resonance imaging or seismic imaging. Our proposed structural sparse point encoding method may also be used in point cloud data analysis. Future work could also explore these potential applications and generalize the work to other research fields.

Acknowledgements

Our work has been funded by NSFC General Research Fund No. 62372393 and No. 12373097.

References

- [1] JG Ables. Maximum entropy spectral analysis. *Astronomy and Astrophysics Supplement*, Vol. 15, p. 383, 15:383, 1974. 1, 3
- [2] Robert D Blum, Kaylan Burleigh, Arjun Dey, David J Schlegel, Aaron M Meisner, Michael Levi, Adam D Myers, Dustin Lang, John Moustakas, Anna Patej, et al. The decam legacy survey. In *American Astronomical Society Meeting Abstracts# 228*, pages 317–01, 2016. 6
- [3] Katherine L Bouman, Michael D Johnson, Daniel Zoran, Vincent L Fish, Sheperd S Doeleman, and William T Freeman. Computational imaging for vlbi image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 913–922, 2016. 1, 3
- [4] Katherine L Bouman, Michael D Johnson, Adrian V Dalca, Andrew A Chael, Freek Roelofs, Sheperd S Doeleman, and William T Freeman. Reconstructing video of time-varying sources from radio interferometric measurements. *IEEE Transactions on Computational Imaging*, 4(4):512–527, 2018. 2
- [5] Andrew A Chael, Michael D Johnson, Katherine L Bouman, Lindy L Blackburn, Kazunori Akiyama, and Ramesh Narayan. Interferometric imaging directly with closure phases and closure amplitudes. *The Astrophysical Journal*, 857(1):23, 2018. 6
- [6] Andrew A Chael, Katherine L Bouman, Michael D Johnson, Ramesh Narayan, Sheperd S Doeleman, John FC Wardle, Lindy L Blackburn, Kazunori Akiyama, Maciek Wielgus, Chi-kwan Chan, et al. ehtim: Imaging, analysis, and simulation software for radio interferometry. *Astrophysics Source Code Library*, pages ascl–1904, 2019. 6
- [7] Shaoyu Chen, Xinggong Wang, Tianheng Cheng, Qian Zhang, Chang Huang, and Wenyu Liu. Polar parametrization for vision-based surround-view 3d detection. *arXiv preprint arXiv:2206.10965*, 2022. 2
- [8] Aleksandra Ćiprijanović, Ashia Lewis, Kevin Pedro, Sandeep Madireddy, Brian Nord, Gabriel N Perdue, and Stefan Wild. Semi-supervised domain adaptation for cross-survey galaxy morphology classification and anomaly detection. *Machine Learning and the Physical Sciences workshop, Advances in Neural Information Processing Systems*, 2022. 6
- [9] The Event Horizon Telescope Collaboration. First sagittarius a* eht results: Calibrated data, 2022. We release a data set to accompany the First Sagittarius A* Event Horizon Telescope Results paper series (EHT Collaboration et al. 2022a,b,c,d,e,f). The data set is derived from the Rev7 Correlation of the Event Horizon Telescope (EHT)’s April 2017 observation campaign (EHT Collaboration et al. 2019c), with further processing and science validation as described in EHT Collaboration et al. 2022b. It is made public simultaneously with four imaging pipelines (2022-D02-02, EHT Collaboration et al. 2022c). This data set contains Sagittarius A* (Sgr A*) data for both low and high bands for two observed days (April 6th and 7th, 2017). Data from the 2017 observations were processed through three independent reduction pipelines (Blackburn et al. 2019, Janssen et al. 2019, M87 Paper III, and Sgr A* Paper II, Paper III). This release includes the fringe fitted, a-priori calibrated, and network calibrated data from both the EHT-HOPS and rPICARD (CASA) pipelines, which are used in the First Sgr A* EHT results. Independent flux calibration is performed based on estimated station sensitivities during the campaign (Issaoun et al. 2017, Janssen et al. 2019, Wielgus et al. 2022). A description of the data properties, their validation, and estimated systematic errors is given in M87 Paper III and Sgr A* Paper II and Paper III with additional details in Wielgus et al. (2019) and Wielgus et al. (2022). The data are time averaged to 10 seconds and frequency averaged over all 32 intermediate frequencies (IFs). All polarization information is explicitly removed. To make the resulting uvfits files compatible with popular very-long-baseline interferometry (VLBI) software packages, the circularly polarized cross-hand visibilities RL and LR are set to zero along with their errors, while parallel-hands RR and LL are both set to an estimated Stokes I value. Measurement errors for RR and LL are each set to $\sqrt{2}$ times the statistical errors for Stokes I. 8
- [10] Liam Connor, Katherine L Bouman, Vikram Ravi, and Gregg Hallinan. Deep radio-interferometric imaging with polish: Dsa-2000 and weak lensing. *Monthly Notices of the Royal Astronomical Society*, 514(2):2614–2626, 2022. 1, 3
- [11] Arjun Dey, David J Schlegel, Dustin Lang, Robert Blum, Kaylan Burleigh, Xiaohui Fan, Joseph R Findlay, Doug Finkbeiner, David Herrera, Stéphanie Juneau, et al. Overview of the desi legacy imaging surveys. *The Astronomical Journal*, 157(5):168, 2019. 6
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceeding of the International Conference on Learning Representations*, 2021. 2, 5
- [13] Leung Henry. Galaxy10 decals dataset. <https://github.com/henrysky/Galaxy10>, 2021. 6, 8
- [14] JA Högbom. Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement Series*, 15:417, 1974. 1, 3, 6, 7
- [15] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010. 6
- [16] Roger A Horn. The hadamard product. In *Proc. Symp. Appl. Math.*, pages 87–169, 1990. 5
- [17] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13919–13929, 2021. 3, 4, 6
- [18] Honghao Liu, Qiong Luo, and Feng Wang. Efficient radio interferometric imaging on the gpu. In *2022 IEEE 18th International Conference on e-Science (e-Science)*, pages 95–104. IEEE, 2022. 3

- [19] Ming Nie, Yujing Xue, Chunwei Wang, Chaoqiang Ye, Hang Xu, Xinge Zhu, Qingqiu Huang, Michael Bi Mi, Xinchao Wang, and Li Zhang. Partner: Level up the polar representation for lidar 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3801–3813, 2023. [2](#)
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [6](#)
- [21] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. [5](#)
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. [2](#), [6](#)
- [23] Kevin Schmidt, Felix Geyer, Stefan Fröse, Paul-Simon Blomenkamp, Marcus Brüggem, Francesco de Gasperin, Dominik Elsässer, and Wolfgang Rhode. Deep learning-based imaging in radio interferometry. *arXiv preprint arXiv:2203.11757*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [24] David R Silva, Robert D Blum, Lori Allen, Arjun Dey, David J Schlegel, Dustin Lang, John Moustakas, Aaron M Meisner, Francisco Valdes, Anna Patej, et al. The mayall z-band legacy survey. In *American Astronomical Society Meeting Abstracts# 228*, pages 317–02, 2016. [6](#)
- [25] Himanshu Singh and Himanshu Singh. Basics of python and scikit image. *Practical Machine Learning and Image Processing: For Facial Recognition, Object Detection, and Pattern Recognition Using Python*, pages 29–61, 2019. [6](#)
- [26] He Sun and Katherine L Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2628–2637, 2021. [1](#), [3](#)
- [27] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. [5](#)
- [28] A Richard Thompson, James M Moran, and George W Swenson. *Interferometry and synthesis in radio astronomy*. Springer Nature, 2017. [2](#), [3](#)
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [2](#), [5](#)
- [30] Ruoqi Wang, Zhuoyang Chen, Qiong Luo, and Feng Wang. A conditional denoising diffusion probabilistic model for radio interferometric image reconstruction. In *26th European Conference on Artificial Intelligence*, pages 2499–2506, 2023. [1](#), [3](#)
- [31] Benjamin Wu, Chao Liu, Benjamin Eckart, and Jan Kautz. Neural interferometry: Image reconstruction from astronomical interferometers using transformer-conditioned neural fields. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [32] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12193–12202, 2020. [2](#)
- [33] Enze Xie, Wenhai Wang, Mingyu Ding, Ruimao Zhang, and Ping Luo. Polarmask++: Enhanced polar representation for single-shot instance segmentation and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5385–5400, 2021. [2](#)
- [34] Yutong Xie and Quanzheng Li. Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VI*, pages 655–664. Springer, 2022. [6](#), [7](#)
- [35] Hu Zou, Xu Zhou, Xiaohui Fan, Tianmeng Zhang, Zhimin Zhou, Jundan Nie, Xiyan Peng, Ian McGreer, Linhua Jiang, Arjun Dey, et al. Project overview of the beijing–arizona sky survey. *Publications of the Astronomical Society of the Pacific*, 129(976):064101, 2017. [6](#)