Astronomy & Astrophysics

Semi-supervised rotation measure deconvolution and its application to MeerKAT observations of galaxy clusters

V. Gustafsson[1],[★] , M. Brüggen[1] , and T. Enßlin[2]

[1] Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany
[2] Max-Planck Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany

ABSTRACT

Context. Faraday rotation contains information about the magnetic field structure along the line of sight and is an important instrument in the study of cosmic magnetism. Traditional Faraday spectrum deconvolution methods such as RMCLEAN face challenges in resolving complex Faraday dispersion functions and handling large datasets.
Aims. We developed a deep learning deconvolution model to enhance the accuracy and efficiency of extracting Faraday rotation measures from radio astronomical data, specifically targeting data from the MeerKAT Galaxy Cluster Legacy Survey (MGCLS).
Methods. We used semi-supervised learning, where the model simultaneously recreates the data and minimizes the difference between the output and the true signal of synthetic data. Performance comparisons with RMCLEAN were conducted on simulated as well as real data for the galaxy cluster Abell 3376.
Results. Our semi-supervised model is able to recover the Faraday dispersion for extended rotation measure (RM) components, while accounting for bandwidth depolarization, resulting in a higher sensitivity for high-RM signals, given the spectral configuration of MGCLS. Applied to observations of Abell 3376, we find detailed magnetic field structures in the radio relics, and several active galactic nuclei. We also applied our model to MeerKAT data of Abell 85, Abell 168, Abell 194, Abell 3186, and Abell 3667.
Conclusions. We have demonstrated the potential of deep learning for improving RM synthesis deconvolution, providing accurate reconstructions at a high computational efficiency. In addition to validating our data against existing polarization maps, we find new and refined features in diffuse sources imaged with MeerKAT.

Key words. magnetic fields – polarization – methods: data analysis – techniques: polarimetric –
galaxies: clusters: intracluster medium

1. Introduction

The utilization of Faraday rotation in astrophysics spans a diverse array of cosmic phenomena, ranging from the interstellar medium of our own Galaxy to extragalactic sources. By studying the amount of Faraday rotation that polarized radio emission suffers from the source to the observer, we learn about the magnetic properties of the intervening medium. The so-called Faraday depth contains information on the strength of the magnetic field along the line of sight. The dispersion of the Faraday depth gives us information about a possible turbulent component of the magnetic field. Modern radio interferometers, such as MeerKAT (Jonas & MeerKAT Team 2016), the LOw Frequency ARray (LOFAR; van Haarlem et al. 2013), and the Australian Square Kilometre Array Pathfinder (ASKAP; Johnston et al. 2007), allow us to study the radio sky across a wide frequency band, while achieving unprecedented precision in polarization measurements.

The MeerKAT radio interferometer (Jonas & MeerKAT Team 2016), situated in South Africa, is particularly suited to polarization studies, owing to its large bandwidth. The dense inner configuration of the MeerKAT array provides a high sensitivity to extended emission, while its longest baseline, spanning 7698 m, offers a high angular resolution. At the central frequency of 1283 MHz, the largest

resolved angular size is about 27.5′ together with a nominal resolution of about 6″.

Multiple methods of decomposing the Faraday rotating signal have been proposed over the years, for example, Faraday rotation measure (RM) synthesis (Burn 1966; Brentjens & de Bruyn 2005; Bell & Enßlin 2012), wavelet decomposition (Frick et al. 2010), compressive sampling (Li et al. 2011; Andrecut et al. 2012), and QU fitting (Farnsworth et al. 2011; O'Sullivan et al. 2012; Ideguchi et al. 2014). For a comparison between the algorithms, readers can refer to Sun et al. (2015). Parametric fits for the spectra of Faraday depth and Stokes Q and U typically assume a Gaussian random distribution for the components of the turbulent magnetic field. Model-free descriptions have been developed, for example, by Van Eck (2018).

In this work we focus on RM synthesis, which decomposes the polarized emission into its constituent Faraday rotating components in a computationally cost-effective way. However, due to the limited coverage in frequency space, uncertainties and false positives in the form of side lobes arise in the Fourier space, particularly for cases with multiple RM components along the line of sight (e.g., Farnsworth et al. 2011; Kumazaki et al. 2014; Miyashita et al. 2016). Methods such as RMCLEAN (Heald et al. 2009) reconstruct the RM synthesis signal iteratively, gradually adding point sources to a model spectrum until the residual is below a specified threshold. While this point source assumption is valid in cases where the emitted

★ Corresponding author;
victor.gustafsson@hs.uni-hamburg.de

synchrotron radiation is separated from the Faraday rotating medium, it breaks down for regions where the radiation is emitted in a magnetized plasma. Examples of this include Galactic diffuse emission, supernova remnants, nearby galaxies, and diffuse radio emission in galaxy clusters (Sun et al. 2015).

Prior works have investigated the application of deep learning techniques to classify the complexity of RM spectra (e.g., Brown et al. 2019; Alger et al. 2021); however, deep learning has not been applied to the deconvolution of RM spectra. In analogy to radio interferometry, where the issue originates from incomplete uv coverage, recent advancements in deep learning have addressed similar issues by methods such as uv-space completion (Schmidt et al. 2022) and direct image reconstruction using diffusion models (Wang et al. 2023). In this study, we propose a novel approach for the deconvolution of RM spectra employing deep learning, by training on data from the MeerKAT Galaxy Cluster Legacy Survey (MGCLS; Knowles et al. 2022), together with simulated data, generated to match the distribution of the observational data.

This paper is organized as follows: In Sect. 2 we describe the basics of RM synthesis. In Sect. 3 we describe the observations used in this work. Section 4 explains the deep learning model, including the data preprocessing, the neural network architecture, and training. In Sect. 5 we present the results from simulated and observational data and compare our results to RMCLEAN. The discussion and conclusions are given in Sects. 6 and 6.4.

2. Faraday rotation measure synthesis

In analogy to radio interferometry, the Faraday dispersion relation can be expressed by the measurement equation, relating the Faraday depth ϕ, to the spectral dimension, defined by the wavelength squared λ^2. Following the notations from Burn (1966) for the Faraday dispersion function $F(\phi)$ and the complex polarized intensity $P(\lambda^2)$, the simplest form of the measurement equation reads

$$P(\lambda^2) = \int_{-\infty}^{\infty} F(\phi)e^{2i\lambda^2\phi}\mathrm{d}\phi + n(\lambda^2),\tag{1}$$

where $n(\lambda^2)$ is noise, which for this work is assumed to be uncorrelated in λ^2 and Gaussian-distributed in the real and imaginary parts. Additional terms, such as channel weights and a channel dependent sensitivity window (Pratley & Johnston-Hollitt 2020) that has been proposed to account for channel-averaging effects, are not used in this work. As we are working with a limited bandwidth, in both λ^2 and ϕ, Eq. (1), reduces to the discrete inverse Fourier transform

$$\tilde{P}(\lambda^2) = \sum_i^N F_i e^{2i\lambda^2\phi_i} + n(\lambda^2),\tag{2}$$

or in matrix form

$$y = \Phi x + n,\tag{3}$$

where Φ is the measurement operator of RM synthesis. Given the increasingly large data sizes generated by modern interferometers, data are often averaged over frequency. While this process increases the signal-to-noise ratio and reduces the computational cost of both calibration and imaging, it also carries the risk of bandwidth depolarization. For Faraday rotating signals, where the real and imaginary parts are a series of sinusoidal waves, the averaging process can smear out the phase information across

Table 1. Spectral and RM synthesis configuration of MGCLS.

Frequency range	856–1712 MHz
– Bandwidth	797 MHz
Number of channels	
– initial	4096
– final	12
Frequency resolution $\delta\nu$	
– initial	209 kHz
– final	62 MHz (mean)
Wavelength squared bandwidth $\Delta\lambda^2$	0.076 m^2
– min resolution $\delta\lambda^2_{\min}$	0.0019 m^2
– max resolution $\delta\lambda^2_{\max}$	0.0168 m^2
Faraday depth bandwidth $\Delta\phi$	512 rad m^{-2}
Faraday depth sampling $\delta\phi$	1 rad m^{-2}

frequency channels, leading to a loss of polarization information. This effect becomes more pronounced for high rotation measures or low frequencies, where the polarization angle rotates rapidly. This can average the resulting signals to near zero. As a result, the maximum Faraday depth to which one has more than 50% sensitivity is approximately $\|\phi_{\max}\| = \sqrt{3}/\delta\lambda^2$.

Here we will write the channel-averaging operator as

$$\mathcal{A}y_i = \frac{1}{N}\sum_j^N y_j,\tag{4}$$

where y_i is the average signal intensity of the i-th channel bin, and N is the number of channels in each bin.

By expressing these operations in matrix form, we can leverage the broadcasting feature of matrix multiplication, enabling efficient element-wise operations across arrays of different shapes without explicit looping. This, coupled with the parallel computing capabilities of GPUs, accelerates the execution of the algorithm for processing large-scale datasets significantly.

3. Observation

The data used in this paper have been taken from MGCLS. For imaging and calibration details see Knowles et al. (2022). Out of the total number of 115 surveyed galaxy clusters, 44 were imaged in full Stokes. The data products include image cubes consisting of 12 frequency channels (908−1656 MHz), within the L-band (856−1712 MHz), at the full 8″ and 15″ resolution to help recover low-surface-brightness features, such as radio relics. Table 1 provides the spectral and RM configuration of this dataset.

The 12 spectral channels of the dataset provide an RM sensitivity up to 173 rad m^{-2}, where the sensitivity is reduced to half. However, as we will see in Sect. 5.1, using deep learning, we can detect higher values of RM by learning the sensitivity window of our observation. Consequently, the selected RM search range is set to $[-256, 256)$ rad m^{-2}, with a sampling in Faraday space of 1 rad m^{-2}. Due to the limited bandwidth of our observation, each point in ϕ will be convolved with the Rotation Measure Spread Function (RMSF), shown in Figure 1, together with an example Faraday spectrum from the Abell 3376 dataset.

The main focus in this paper is the Abell 3376 cluster, which has previously been studied by, for example, Kale et al. (2012), George et al. (2015), and Chibueze et al. (2023). The studies mostly focused on the radio relics. In Hu et al. (2024), synchrotron intensity gradient (SIG) mapping was used to infer
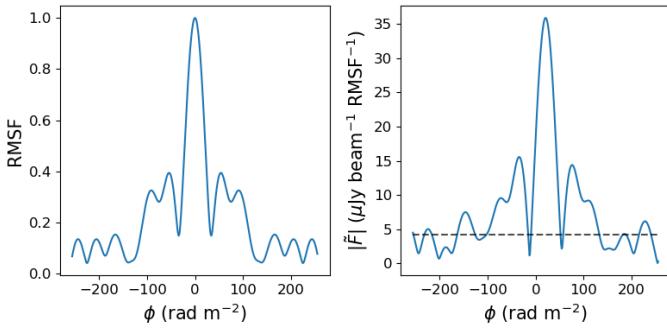
Fig. 1. Magnitude of the RMSF for the MGCLS band (left) and an example Faraday spectrum taken from a pixel in the eastern relic (right). The estimated noise level is shown as a dashed black line.

the magnetic field orientation in the radio relics. However, no studies have yet been conducted with MeerKAT in full polarization.

This cluster harbors two radio relics that extend for Megaparsecs. No radio halo has been detected in this cluster. The orientation of the cluster relics, together with the extended X-ray emission stretching in the northwest-southeast direction (Kale et al. 2012), suggests that Abell 3376 is a merging cluster.

Furthermore, the cluster contains several radio galaxies and active galactic nuclei (AGN). Most notably close to the eastern relic, with jets bent by ~90° from their original direction, which is suggested to be caused by the cluster magnetic field (Chibueze et al. 2021).

In Appendix A we briefly go over the results of applying the deep learning model to five more datasets from the MGCLS DR1, and compare them with previous studies.

4. Deep learning deconvolution

The neural network model aims to invert Eq. (3), which generally lacks invertibility. However, we can approximate a highly nonlinear function tailored to this specific task by fine-tuning the parameters of a neural network. The parameter optimization process uses the ADAM optimizer (Kingma & Ba 2017), aiming to minimize the mean squared error (MSE) loss between the observed data and the result obtained after applying the measurement and channel-averaging operators to the network's predictions, as given by:

$$L = \frac{1}{N} \sum_i^N \|y_i - \mathcal{A}\Phi\hat{x}_i\|^2, \tag{5}$$

where N is the number of samples in each batch, $\hat{x}_i$ are the model outputs in ϕ and y_i are the data in λ^2. The motivation behind using the MSE loss lies in the assumption of Gaussian-distributed noise in both the real and imaginary parts of the complex polarized intensity. This choice is based on the observation that the log likelihood function for a Gaussian distribution is directly linked to the squared Euclidean norm, as seen in the MSE formulation. This approach ensures that our loss function captures the statistical characteristics of the data while minimizing the error between predicted and observed values.

Minimizing Eq. (5) allows us to identify a solution that aligns with the observed data. However, given the problem's infinite solution space, we require a prior to guide the optimization toward a physically realistic solution. Previous work, such as

nonparametric QU fitting Pratley et al. (2021), incorporate an ℓ_1 regularization term to constrain the number of RM components, essentially adopting a CLEAN prior approach. Instead, this work includes a Faraday-thick prior by utilizing semi-supervised learning, including samples with simulated Gaussian sources in ϕ, in the training dataset, and allowing the model access to the true signal of those samples. The model should thus find a general solution, that reconstructs the simulated sources, while also fitting the observed data. The full semi-supervised loss thus reads

$$L = \frac{\beta}{N_\text{obs}} \sum_i^{N_\text{obs}} \|y_i - \mathcal{A}\Phi\hat{x}_i\|^2 + \frac{1}{N_\text{sim}} \sum_j^{N_\text{sim}} \|x_j - \hat{x}_j\|^2, \tag{6}$$

where the loss is divided into contributions from real samples (N_obs) and simulated samples (N_sim). The first term of the loss function thus compares the spectra in λ^2, while the second part compares the output $\hat{x}_j$ with the true signal x_j in ϕ. The factor β is used for weighting the different loss terms, assuring that information is passed from the supervised learning onto the observational data. The β value was set to balance the two loss terms so that they were of the same order of magnitude. Since the signal in λ^2 spans the entire bandwidth, unlike the sparse peaks seen in ϕ, the first loss term would dominate without appropriate weighting. Thus, a value of 10^{-5} was found to be effective, though this value was not extensively optimized.

4.1. Preprocessing

Creating a balanced dataset of synthetic data is a relatively simple process. However, observational data are typically unbalanced and may include outliers, which can negatively affect the training process. Among the samples, bright point sources have the most significant impact. Additionally, to avoid training the model on data predominantly comprised of noise, noise dominated samples were also excluded from the training dataset. As the polarized intensity map was significantly contaminated by foreground emission, a threshold was instead set by the noise level in total intensity σ_I, calculated as

$$\sigma_I^2 = \frac{1}{N} \sum_i^N \sigma_{I,i}^2, \tag{7}$$

where N is the number of channels in the Stokes I cube, and $\sigma_{I,i}$ is the noise in the i-th channel, measured from an emission-free region. The implications of setting the threshold in total intensity is discussed in Sect. 6.2. Both noisy samples and bright sources were excluded by clipping the dataset to within the range of $(3, 30)\sigma_I$. The full preprocessing and training steps are shown in Fig. 2.

After sampling the observed data, $\tilde{P}_\text{obs}(\lambda^2)$, we extracted the peak flux A, RM, and intrinsic polarization angle θ from the RM synthesis signals $\tilde{F}_\text{obs}(\phi)$. The full set of simulation parameters can be seen in Table 2. These parameters were then used to create one-dimensional simulated samples, each consisting of a Gaussian mixture model, with one to five Gaussian components in ϕ. We note that these simulations are a significant simplification of the true profile of Faraday spectra, and that realistic Faraday spectra are expected to exhibit non-Gaussian features. Examples of such features can be seen in Bell et al. (2011), where sharply peaked and asymmetric profiles are obtained from more advanced simulations.

The amplitude of each component was adjusted to conserve the total flux of the sample by dividing it by the number of
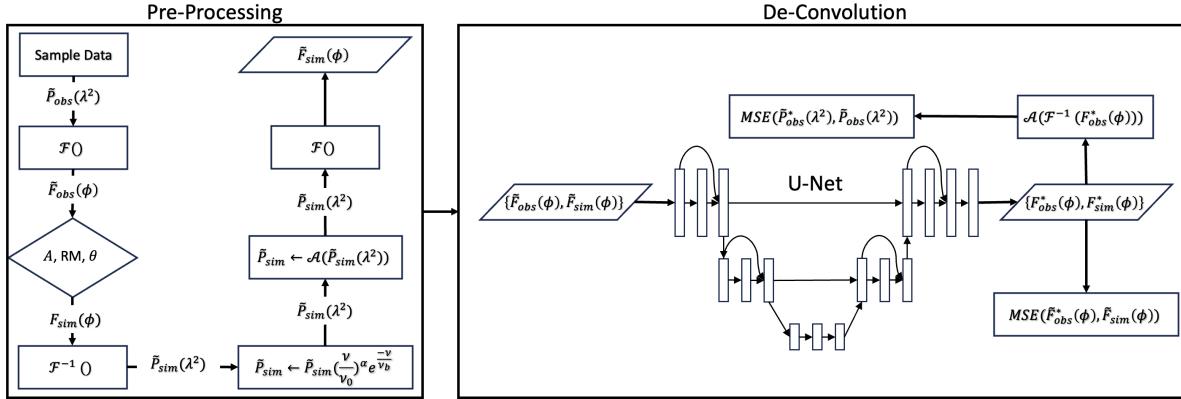
Fig. 2. Flowchart illustrating our deep learning deconvolution model. The notations $\mathcal{F}$ and $\mathcal{F}^{-1}$ signify the transformations from λ^2 to ϕ and back, respectively. Initially, samples are extracted from the observed data and used to generate a simulated dataset with a similar distribution. Next, the two datasets are combined and fed into the neural network. After the deconvolution process, the predictions F^*_{obs} and F^*_{sim} are separated once more and evaluated against their corresponding targets using the MSE loss.

Table 2. Source parameters for creating a simulated dataset that matches the data distribution.

Number of sources	$U(1, 5)$
Amplitude A (µJy)	$\sim\mathcal{D}$
Rotation measure RM (rad m^{-2})	$\sim\mathcal{D}$
Intrinsic polarization angle θ (radians)	$\sim\mathcal{D}$
FWHM (rad m^{-2})	$U(1, 10)$
Spectral index α	$U(-3, 0)$
Breaking frequency ν_{b} (MHz)	$U(1, 10)$

Notes. The amplitude A, RM, intrinsic polarization angle θ are drawn from the data $\mathcal{D}$. The number of sources, full width at half max FWHM, spectral index α, and breaking frequency are drawn from uniform distributions.

Table 3. Hyperparameters for the deep learning deconvolution model.

Number of levels	5
Minimum number of channels	64
Convolution kernel size	3
Dropout rate	0.2
Batch size	256
Initial learning rate	10^{-5}
– min learning rate	10^{-7}
– learning rate factor	0.1
– learning rate patience	5 epochs
Loss weighting factor β	10^{-5}

components. Since the peak flux of the RM synthesis signal corresponds to the integral of an unresolved source, the amplitude was also adjusted by dividing it by the factor that relates the area of a Gaussian to its amplitude, given by $2\sqrt{\log(2)}/\sqrt{\pi}$ FWHM. The full width at half maximum (FWHM) of each Gaussian component was drawn from a uniform distribution between 1 and 10 rad m^{-2}. Each simulated sample was then transformed to λ^2, where each signal was multiplied by a combination of a power-law and an exponential:

$$\tilde{P} = \tilde{P}_0 \left(\frac{\nu}{\nu_0}\right)^\alpha e^{-\nu/\nu_{\mathrm{b}}}, \tag{8}$$

where α is the spectral index, ν_0 is set to the lowest frequency in our band, 856 MHz, and ν_{b} is the breaking frequency. No noise is added at this stage for reasons discussed in Section 6.2. Samples were then averaged to the 12 frequency channels provided in the dataset. The simulated signals were then transformed to ϕ, resulting in $\tilde{F}_{\mathrm{sim}}(\phi)$ and paired with their corresponding true signal $F_{\mathrm{sim}}(\phi)$ to create input-target pairs. Similarly, each observed sample $\tilde{F}_{\mathrm{obs}}(\phi)$ was paired with its respective λ^2 spectrum $\tilde{P}_{\mathrm{obs}}(\lambda^2)$ to create input-target pairs for the observation dataset.

4.2. Network architecture

The network architecture for this work is a one-dimensional U-Net (Ronneberger et al. 2015). The network consists of five

levels, each consisting of a ResNet block (He et al. 2015). In the contracting path, max-pooling is used to progressively down-sample the input by a factor of two after each ResNet block, while in the expanding path, nearest neighbor up-sampling is applied. Additionally, the model employs skip connections with attention gates (Oktay et al. 2018) to retain fine structural details. At the bottleneck, the downsampled input should contain most of the signal information in a compact latent space. From this high-dimensional representation, the model should reconstruct the signal, effectively performing the deconvolution process and recovering the true Faraday dispersion function. In Table 3, the set of hyperparameters used for deconvolving the Abell 3376 dataset are shown. The model uses a learning rate scheduler, monitoring Eq. (6) during training and decreasing the learning rate by a factor if the loss is not reduced after a set number of epochs. This approach allows for starting with a relatively high learning rate, thus speeding up convergence, while also reducing the risk of overfitting.

4.3. Training

The model was trained with 80 000 real and 20 000 simulated samples over 100 epochs, reaching convergence after approximately 80 epochs. The dataset was divided into 80% training, 10% validation, and 10% testing data. During training, the model optimizes the network parameters using the training data. The validation dataset is used to monitor the model's performance and tune hyperparameters, preventing overfitting by ensuring
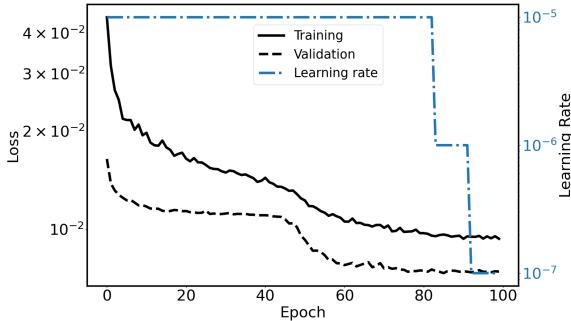
Fig. 3. Loss curves and learning rate during training on the Abell 3376 dataset. The plot shows the training loss (solid line) and validation loss (dashed line) as a function of epoch, plotted on the left y-axis with a logarithmic scale. The learning rate, plotted on the right y-axis with a logarithmic scale, is also shown.

that the model generalizes well to unseen[1] data. The final evaluation of the model's performance is conducted on the test dataset, which provides an unbiased assessment of its predictive accuracy. In Fig. 3, Eq. (6) for the training and validation set is plotted as a function of epochs during the training process, together with the learning rate. We see that the model generalizes well to unseen data, as the validation loss is always lower than the training loss. This is due to the use of dropout, which randomly sets a fraction of the neurons to zero during training, in order to find a simpler and more general solution to the problem. We observe that after 40 epochs, the model encounters a local minimum, where the loss reduction slows down. However, the model manages to escape this minimum and continues to reduce the loss after approximately ten more epochs. At around 80 epochs, the model stops reducing the loss further, prompting the learning rate scheduler to lower the learning rate.

5. Results

In this section we present the results from applying our semi-supervised model first to a set of simulated data, and second to a MeerKAT observation of the galaxy cluster Abell 3376.

5.1. Simple model of Gaussian components

In order to evaluate the recovery of the true signal, we apply our model to a dataset of simulated Gaussian sources. The simulations were created similarly to Sect. 4.1, but with the parameters in Table 4. The noise level σ_{QU} is added to the real and imaginary parts in λ^2, set to the square root of the number of channels, resulting in a noise level of 1 μJy after averaging over all channels.

In Figs. 4 and 5, we compare the performance of the deep learning model and RMCLEAN. A fairer comparison would be to use a multi-scale version of the CLEAN algorithm, similar to Cornwell (2008), which adds components of gradually increasing scale sizes to the model image. As no such implementation exists for RMCLEAN, we should not compare the deep learning model output with the clean components directly, but rather look at peak flux, RM and polarization angle of the low-resolution spectra.

Fig. 4 presents a case with a single source located at $\phi = -85\,\text{rad}\,\text{m}^{-2}$. Our model successfully recovers the signal and

---
[1] Commonly used word in machine learning, data that the model has not optimized its parameters to.

Table 4. Source parameters for a simulated testing dataset.

Number of sources	$U(1, 3)$
Amplitude A (μJy)	$U(0, 50)$
Rotation measure RM (rad m^{-2})	$U(-200, 200)$
Intrinsic polarization angle θ (radians)	$U(-\pi/2, \pi/2)$
FWHM (rad m^{-2})	$U(1, 10)$
Spectral index α	$U(-3, 0)$
Breaking frequency (GHz) ν_b	$U(1, 10)$
σ_{QU} (μJy)	$\sqrt{N_{\text{channels}}}$

Notes. All parameters are drawn from uniform distribution except the noise level σ_{QU} which is the same for all samples.

aligns well with the data in λ^2. While there are minor differences between the true source and the output, they are reduced when the resolution is set to the theoretical limit.

RMCLEAN iteratively adds point sources until a specified threshold is reached, usually stopping just before the model starts cleaning the noise. In our specific example shown in Fig. 4, RMCLEAN places point sources at locations of strong emission. However, rather than precisely recovering the true signal, it overestimates the flux at these points, leaving regions near the peak flux blank. To replicate the Faraday thick component in this scenario, RMCLEAN also places sources further from the peak, leading to a spread that does not accurately reflect the true distribution. While this behavior is observed in our case, especially with Faraday thick components, it should be emphasized that for Faraday thin components, RMCLEAN can more accurately recover the true signal. The performance of RMCLEAN largely depends on the specific characteristics of the signal and how the algorithm is applied.

A more complex Faraday spectrum is shown in Fig. 5, with three sources with different parameters as listed in Table 4. While the deep learning model accurately recovers the strong sources, it fails to detect the weaker source which lies close to the noise level. For the complex spectrum, it becomes even more evident that the RMCLEAN assumption is not valid for extended emission in ϕ.

Figure 6 shows the predicted Gaussian parameters from the reconstructed Faraday dispersion, where each sample includes a single Gaussian component with parameters specified in Table 4. Each detected signal is fitted with a Gaussian model, from which the amplitude, FWHM, and peak RM are extracted. Since the fit is computed using the least squares method, we cannot directly fit a complex function. Hence, the complex phase is determined from the obtained peak RM, using

$$\theta = \frac{1}{2}\arctan\frac{\mathfrak{I}(F(\text{RM}))}{\mathfrak{R}(F(\text{RM}))}, \tag{9}$$

where $\mathfrak{R}$ and $\mathfrak{I}$ are operators that select the real and imaginary parts, respectively.

For all parameters we obtain a clear correlation, but with some having significant errorbars. We note that the amplitude does not follow $U(0, 50)$, due to the spectral scaling as shown in Eq. (8). As a result, the model struggled to accurately predict high amplitudes, which are infrequent in the data. For the FWHM, we observe a clear correlation down to 2 rad m^{-2}, which is unexpected given that the theoretical resolution, defined by the FWHM of the RMSF, is 43 rad m^{-2}. For the RM and intrinsic polarization angle, the errors are minimal, with only about 1.2% of the predicted polarization angles deviating by more than 5°. Interestingly, the largest errors occur for high signal to noise
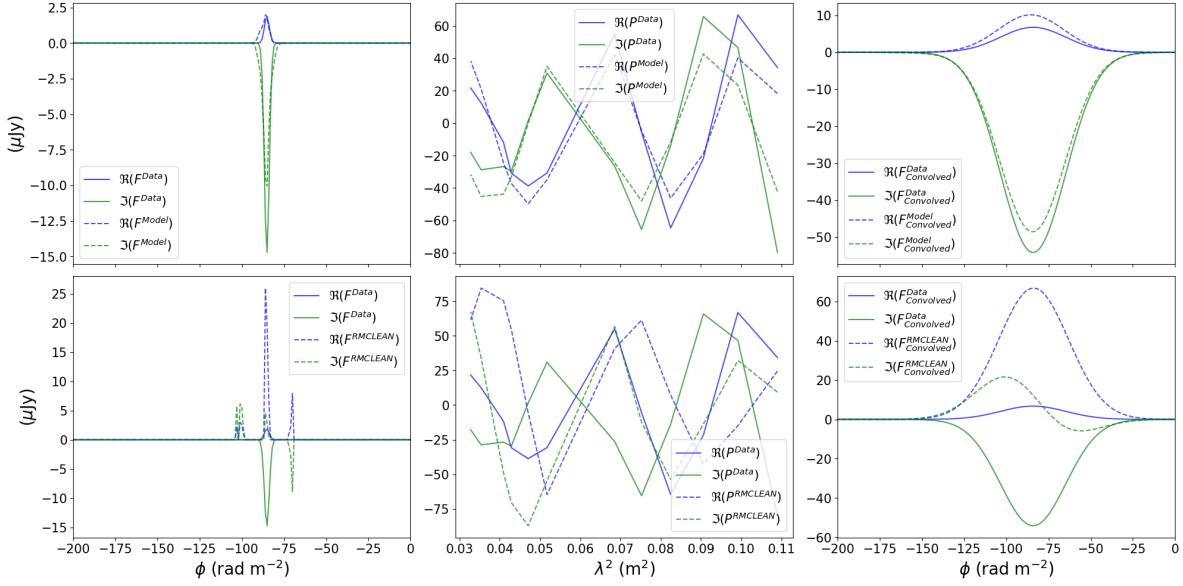
Fig. 4. Comparison between deep learning model output (top) and RMCLEAN (bottom). The outputs are shown in ϕ (left) and λ^2 domain (middle). In the right plot the outputs are tapered to the theoretical resolution of $43\,\text{rad\,m}^{-2}$.
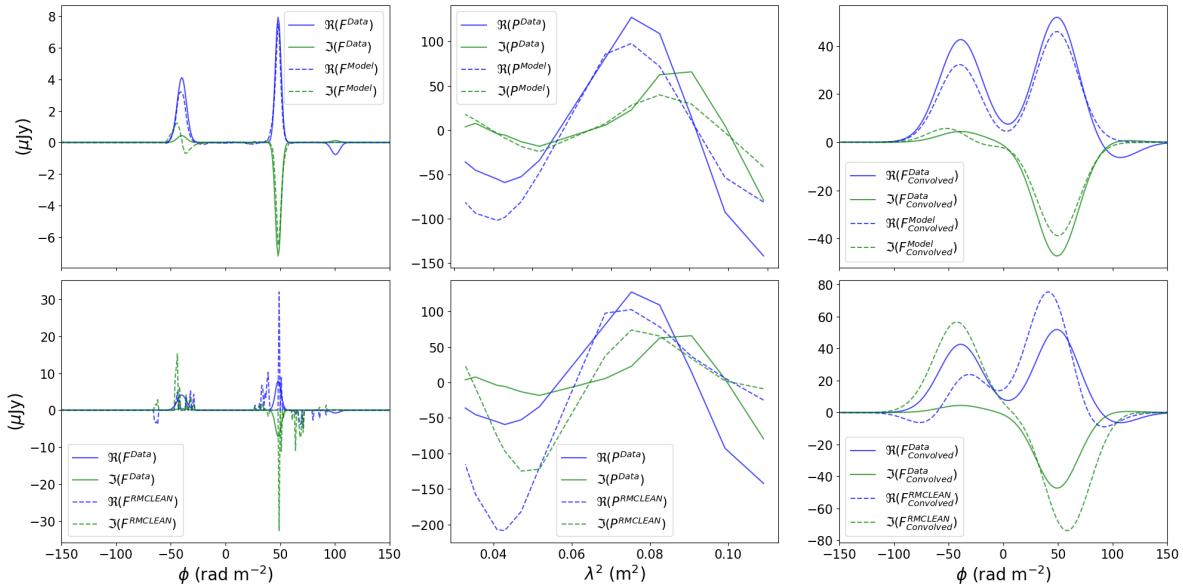

Fig. 5. As Fig. 4, but with an assumed more complex Faraday spectrum.

(S/N) signals, likely due to the dataset imbalance, which causes the model to perform better on the more common low S/N cases.

Since RMCLEAN focuses on one RM component at a time, it cannot account for interactions between overlapping components after RM synthesis. This limitation also applies to extended RM components, which can be viewed as a collection of signals that are closely spaced in ϕ. In contrast, the advantage of the deep learning model is that it considers the entire spectrum simultaneously, allowing it to learn and account for potential interference between signals. This capability is shown in Figure 7, in which we compare the predicted to the true intrinsic polarization angle for samples containing multiple RM components, according to Table 4.

While the errors are larger than for a single Gaussian component, we observe that the predictions are centered on the true polarization angle. We also note that there is a larger spread when the samples include multiple sources, as expected due to greater interference between signals. Quantitatively, the mean absolute errors are 3.7, 13, and 21° for samples containing one, two, and three sources, respectively. As the number of sources increases, the complexity of disentangling overlapping polarization signals rises, leading to higher errors. However, due to the inherent nonlinearity of the neural network, averaging nearby pixels, as demonstrated in Sect. 6.2, is expected to reduce errors, even when individual samples exhibit some similarity. This process leads to smoother variations that more accurately trace the magnetic field lines. Nevertheless, this effect is not explored further in this paper, and a more detailed analysis would be necessary to quantify the extent of this error reduction and identify the conditions under which it is most effective. In contrast, as demonstrated in Miyashita et al. (2016), the predicted polarization angles obtained by reconstructing interfering RM
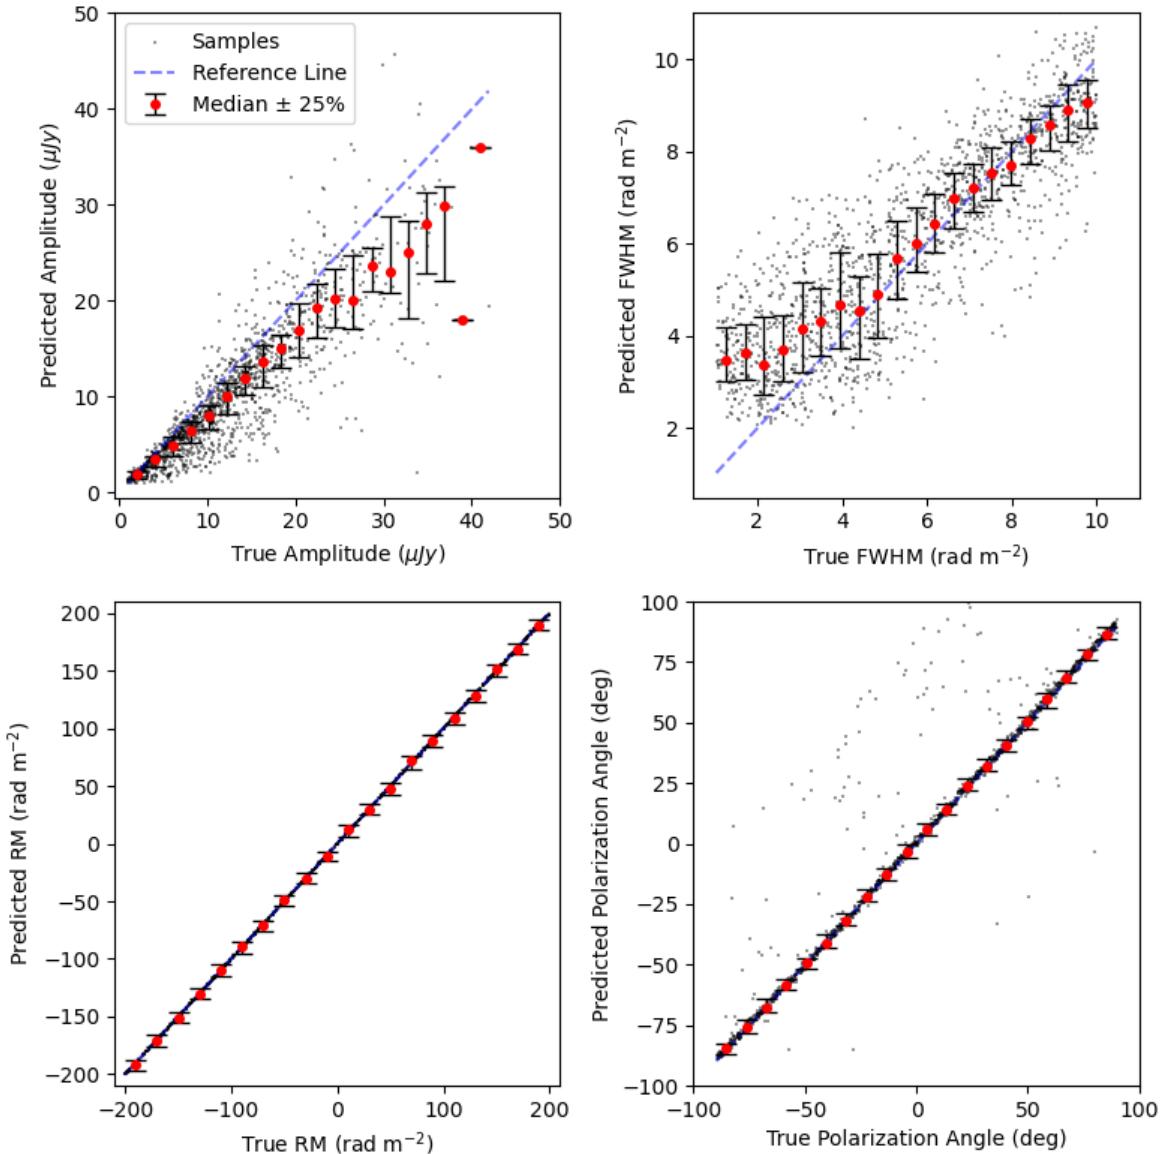
Fig. 6. Predicted versus true parameters from Gaussian fittings of the Faraday dispersion. The parameters are amplitude (top left), FWHM (top right) RM (bottom left) and intrinsic polarization angle (bottom right). Predicted values are shown as black dots. The range is divided into 20 bins from which the median ±25% are calculated and shown as red dots and errorbars respectively. A reference line where the predicted value equals the true value is shown as a dashed blue line. Predictions of the polarization angle with an error larger than $\pi/2$ radians are adjusted by $\pm\pi$ to account for the π-periodicity of the signal.

components with RMCLEAN exhibit systematic errors that do not decrease when averaged over multiple pixels.

As mentioned in Sect. 2, the sensitivity of RM synthesis is dependent on the channel width and the RM of the source, as rapidly rotating signals will lead to small net signals in broad channels. In Fig. 8 we compare the sensitivity of RMCLEAN and the deep learning model. The test is conducted with a single Gaussian with the rest of the parameters according to Table 4, but with RMs up to $\pm512\,\mathrm{rad\,m^{-2}}$. While RMCLEAN demonstrates high accuracy at lower RMs, we observe that its sensitivity is significantly dependent on the RM value. After $\pm400\,\mathrm{rad\,m^{-2}}$, almost no signals are detected. The deep learning model has a lower accuracy at low RMs, but the sensitivity remains almost constant over the entire search range, reaching RMs up to $512\,\mathrm{rad\,m^{-2}}$. As the model is given access to the intrinsic intensity of the source, it implicitly learns the sensitivity

window of the observation, and adjusts the output accordingly. In total, the deep learning model and RMCLEAN are able to find 88% and 72%, respectively, out of the simulated signals.

5.2. MeerKAT data of Abell 3376

Next, we proceeded to test our model on the Abell 3376 dataset from the MGCLS. RMCLEAN was run to a level of $3\sigma_P$ ($\sigma_P = 4.3\,\mu\mathrm{Jy}$) with a gain of 0.10. Observe that $3\sigma_P$ is a rough estimate, computed from regions that contains foreground emission, resulting in an overestimate of the noise level. To have a fair comparison between the models, both algorithms were run without masking any pixels. The deep learning model was run on a NVIDIA A100 80 GB GPU. The computation times for training and inference were 10 and 15 minutes, respectively, while RMCLEAN took 5 h on a single thread of a 64-core AMD
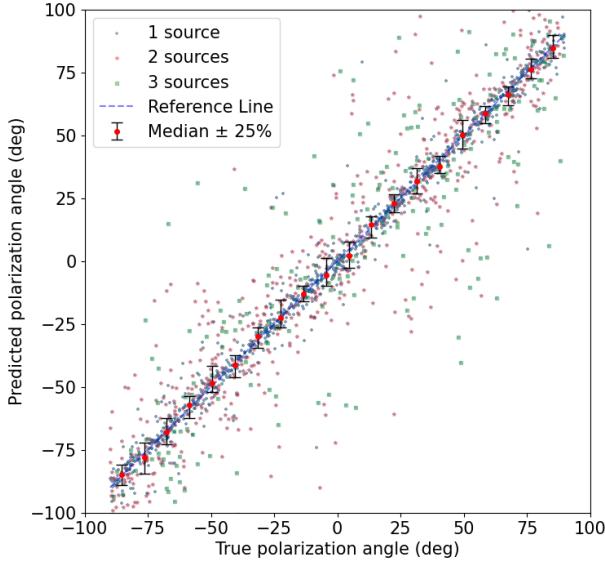
Fig. 7. Predicted versus true intrinsic polarization angle for a test with multiple Gaussian components generated according to Table 4. Predicted angles are shown as black, red and green markers for samples containing one, two and three peaks respectively. The range is divided into 20 bins from which the median ±25% are calculated for the full dataset, and shown as red dots and errorbars respectively. A reference line where the predicted angle equals the true angle is shown as a dashed blue line. Predictions with an error larger than $\pi/2$ radians are adjusted by $\pm\pi$ to account for the π-periodicity of the signal.

Fig. 8. Sensitivity comparison of the deep learning model and RMCLEAN. The range is divided into 20 bins from which the median ±25% are calculated and shown as dots and errorbars, respectively.

EPYC 7H12 CPU, running at a clock frequency of 2.6 GHz. As the current version of RMCLEAN does not support multi-threading, and only considers peaks above $3\sigma_P$ for cleaning, we cannot compare the computation times between the two algorithms.

In Fig. 9 we show how the Faraday dispersion varies in the relics and the bright AGNs. In the plot seven spectra are shown. Five come from line-of-sights through the radio relics in the cluster periphery, and two spectra from the bright AGNs close to the northeastern relic. We observe that where there is strong emission, the RM synthesis signals are deconvolved into narrow peaks, indicating that the magnetic field remains uniform at the scale of the beam. In contrast, at points of weaker emission, typically closer to the cluster center, the Faraday spectrum is more complex. It is worth noting that the polarization cubes were contaminated by foreground emission, which dominates the regions beyond the cluster relics. Therefore, signals in these areas should be interpreted with caution. In the AGN, we observe a clear double peak signal, separated by about 150 rad m^{-2}.

In Fig. 10 the orientation of the polarization vectors, rotated by 90° to show the projected magnetic field orientation is shown. The image is produced after deconvolving the RM synthesis spectra according to the procedure in Sect. 4 and extracting the intrinsic polarization angle of the polarized emission from the highest peak in ϕ. We observe that the magnetic field lines along the outer edge of the cluster relics align with their shapes, particularly evident in the tail of the northeastern relic. This supports the theory that relics are a tracer of merging events, as the compression at the shock front is expected to align the magnetic field with it (e.g., Enßlin et al. 1998; Domínguez-Fernández et al. 2021). As we move inward along the merger axis, the orientation becomes more turbulent and in the regions of weak emission, the orientation is seen as mostly random.
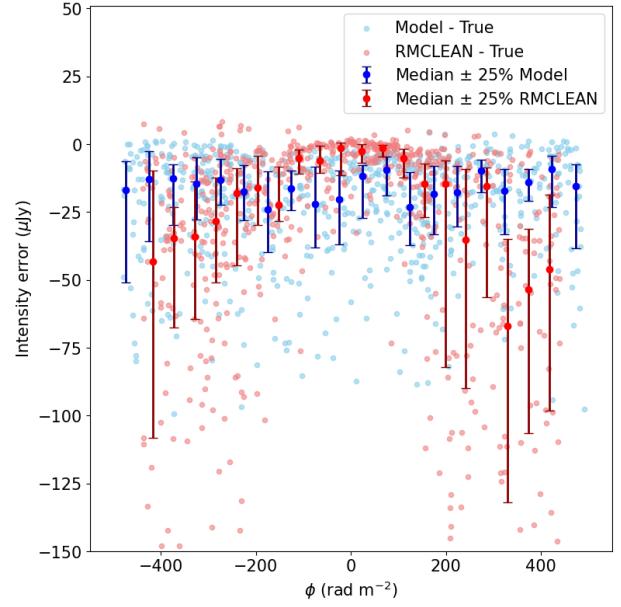
We observe that in most regions, the magnetic field orientation aligns with the dominant filament. However, there are some deviations from this trend, most notably at the upper part of the eastern relic tail (RA = 6^{h}03^{m}00^s, Dec = −40°00′00″), where the magnetic field is oriented toward the inner part of the cluster. This is one of the discrepancies with Hu et al. (2024), where the orientation is perpendicular to the intensity gradient.

In Fig. 11, we present a close-up view of the radio galaxy MRC 0600−399 ($z = 0.04559$). The image reveals that despite the jets bending from their initial trajectory, they stay collimated for approximately 100 kpc in the northern jet and 50 kpc in the southern jet beyond the bend point. Additionally, to the east of the galaxy, we observe a distinct structure, which, based on its higher redshift ($z = 0.0480$), has previously been identified as a separate galaxy (Chibueze et al. 2021). Furthermore, we observe that in both galaxies, the magnetic field orientation far from the AGN aligns with the direction of the jet outflow. This alignment is expected, as the jets can influence the large-scale magnetic field structure by dragging or stretching the field lines along their paths. However, in regions closer to the AGN, where the environment becomes more turbulent due to increased activity and chaotic processes near the black hole, the magnetic field no longer follows this alignment. In some cases, the field lines appear perpendicular to the jet outflow. This deviation is likely due to the complex interactions between the intense radiation pressure, plasma dynamics, and rotational motion in the vicinity of the AGN core, which disrupt the ordered structure of the magnetic field.

In Fig. 12 we show the AGN connected to the eastern relic. This AGN has been proposed by Chibueze et al. (2023), to provide the relic with seed electrons that potentially are reaccelerated by the shock. The link between the AGN and relic is confirmed by the magnetic field connecting the two structures. Additionally, there is a bright spot northwest of the protruding AGN. While the magnetic field in this region aligns with that of the relic, the Faraday spectrum reveals a complex source
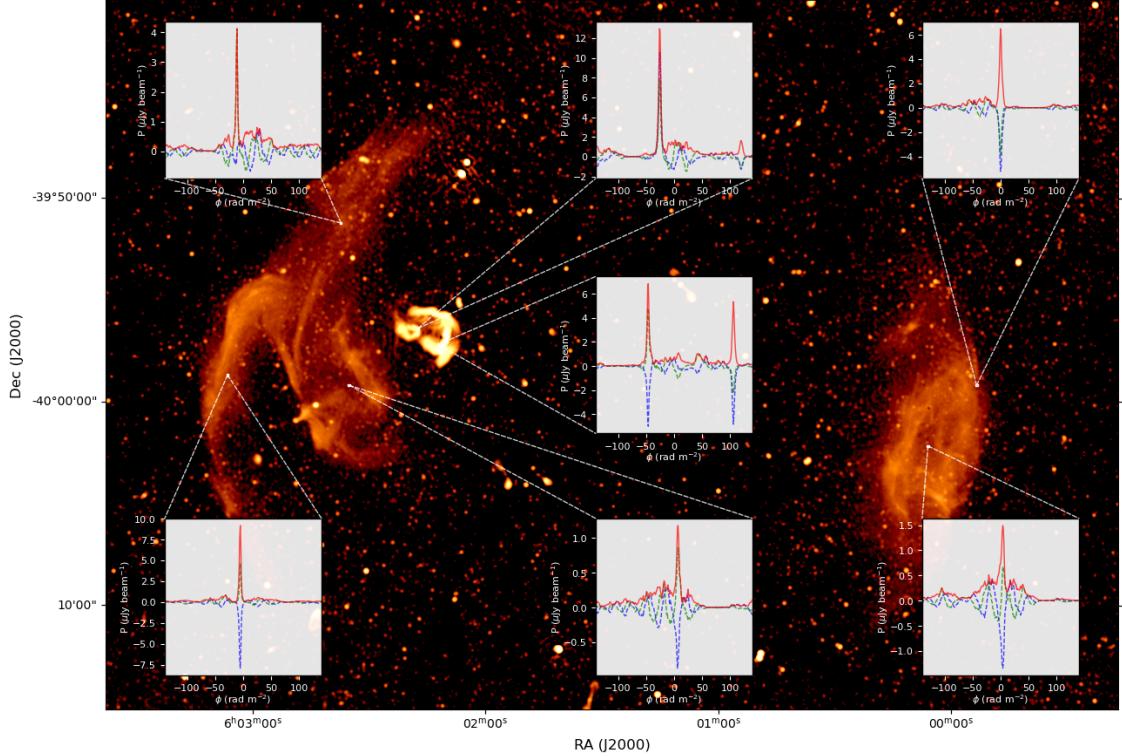
Fig. 9. Total intensity map of Abell 3376 at a resolution of $8'' \times 8''$, with five example spectra from the cluster relics and two from the bright AGNs. Each spectrum has been shifted by local estimates of the Galactic Faraday rotation. The spectra are color-coded as in Figs. 4 and 5, together with the magnitude shown in red.
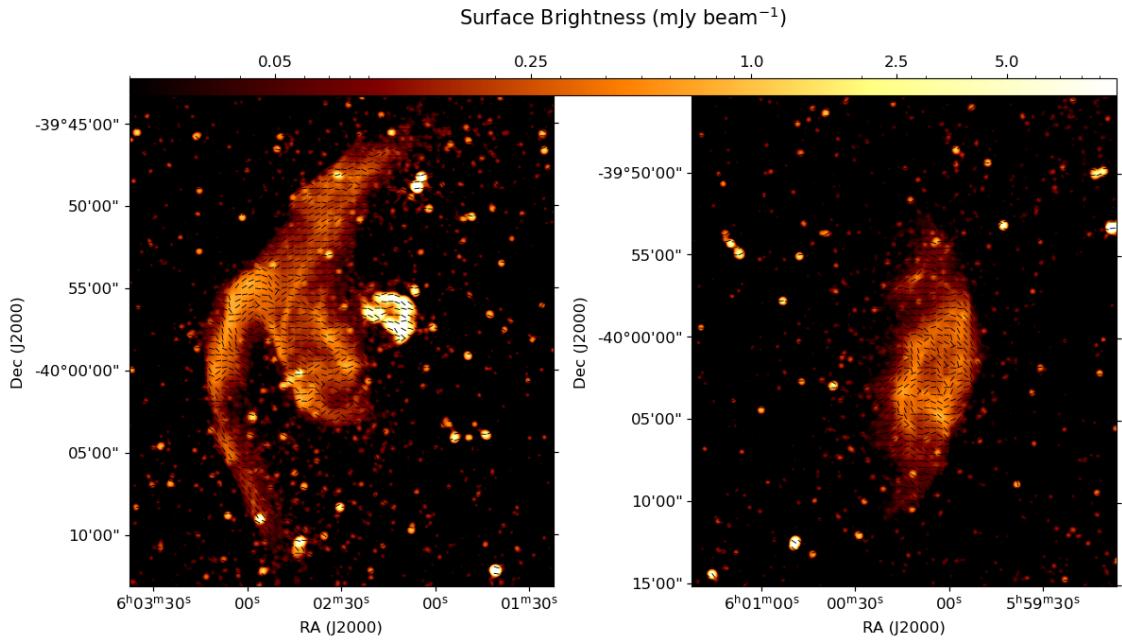

Fig. 10. Total intensity map of the relics in Abell 3376 at a resolution of $15'' \times 15''$. Overlaid are the projected magnetic field vectors.

extended in ϕ, in contrast to the narrow, well-defined peak in the surrounding region.

The rotation measure map of Abell 3376 is shown in Fig. 13. The image is produced by localizing the highest peak in ϕ for each pixel, after which the galactic RM, taken from Hutschenreuter et al. (2022), is subtracted. For the most part of the relics, the RM appear to vary smoothly, except for the weak emission regions, where the peak emission RM values become very unstable, possibly indicating the sensitivity limit for the network. Furthermore, the RMs in the AGN show strong fluctuations, reaching values of up to $\pm 100\,\mathrm{rad\,m^{-2}}$.

In Fig. 14 we show the fractional polarization map of Abell 3376. The map is produced from the Stokes I, Q, U cubes before RM synthesis, by summing along the frequency axis as

Surface Brightness (mJy beam^{-1})
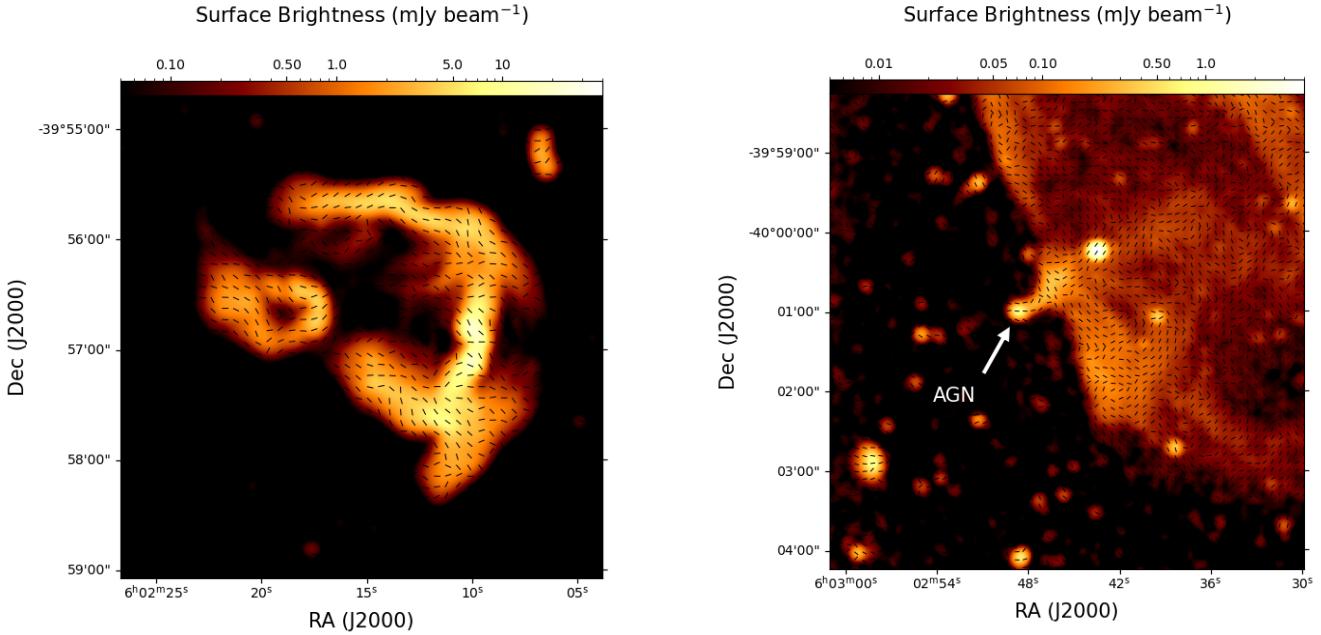
Fig. 11. Total intensity map of the AGNs west of the eastern relic, at a resolution of $8'' \times 8''$. Overlaid are the projected magnetic field vectors.

Surface Brightness (mJy beam^{-1})

Fig. 12. Total intensity map of the AGN connected to the eastern relic, at a resolution of $8'' \times 8''$. Overlaid are the projected magnetic field vectors.

$$\frac{|P|}{I} = \sum_i \frac{\sqrt{Q_i^2 + U_i^2}}{I_i}. \tag{10}$$

Due to the difficulty in accurately computing σ_P, the fractional polarization is not corrected for Rician bias. Therefore, the reported fractional polarization values should be considered as upper limits. At the shock front of the southwestern relic, where the magnetic field aligns with the shock, we observe high fractional polarization, reaching up to 60%. Similarly, in the tail of the northeastern relic, the fractional polarization also reaches high values, up to 50%. While such values are expected in regions with ordered magnetic fields, we also observe peaks in the polarization fraction at the inner edge of the relics, where the magnetic field is not ordered on large scales. These peaks are likely not physical and may instead result from foreground emission, which does not appear in the total intensity map.

6. Discussion and conclusions

In this section we discuss our results on, both, the simulated and observational data. We also discuss some implications of our deep learning model.

6.1. Simulations

In the simulations presented in Section 5.1, we assumed a constant intrinsic polarization angle over each Gaussian component. Thus, for individual sources, the projected magnetic field would be aligned at each point of emission along the line of sight. Although this simplification might not fully capture the physical reality, modeling the intrinsic polarization variations along the line of sight is challenging. Furthermore, due to the limited spatial resolution, each pixel also contains information from neighboring pixels. In the case of a Faraday thick component with

ordered magnetic fields, we would expect some kind of correlation of the polarization angle along ϕ, while in the case of turbulent fields, the polarization angle can be completely random. It is crucial to note that the simulation outputs must still fit the observational data. If the initial polarization angle assumption is inaccurate, the model will adjust it accordingly.

While RMCLEAN is able to locate the peak position of the spectra seen in Figs. 4 and 5, it does not extract the intrinsic polarization angle in a correct manner for a scenario where the polarized intensity follows extended Gaussian profiles. As the algorithm extracts the real and imaginary parts of the Faraday dispersion function from the rapidly varying RM synthesis signal, the polarization angle is not expected to be correct outside of the peak location. This is seen in the secondary peaks in Figs. 4 and 5, where the real and imaginary parts do not correspond to the true signal. Furthermore, RMCLEAN uses the shift theorem of Fourier theory to rotate the RM synthesis signal to the average λ^2. While no information is lost in this transform (Rudnick & Cotton 2023), it means that a small error in the peak position, will cause a significant error of the polarization angle when rotated back to $\lambda^2 = 0$ (Brentjens & de Bruyn 2005).

Fig. 6 shows that the model systematically underestimates the signal amplitude, particularly for samples with high S/N. This is likely because the dataset is unbalanced with respect to signal strength, causing the model to be biased toward lower amplitudes. To resolve this issue, one thus have to make sure that the dataset is balanced, after applying Eq. (8), however, this is left for future work.

From Fig. 6 we saw that the low resolution in Faraday depth, ϕ, can be improved by deep learning, at least for a single unresolved source. While there is no way of confirming the super-resolution on real data, where the true signal is unknown, we can enhance our confidence in the model's capabilities by using neighboring pixels as additional information. Since the model takes one-dimensional inputs and does not account for neighboring pixels, incorporating this spatial domain information should make the super-resolution results more robust.
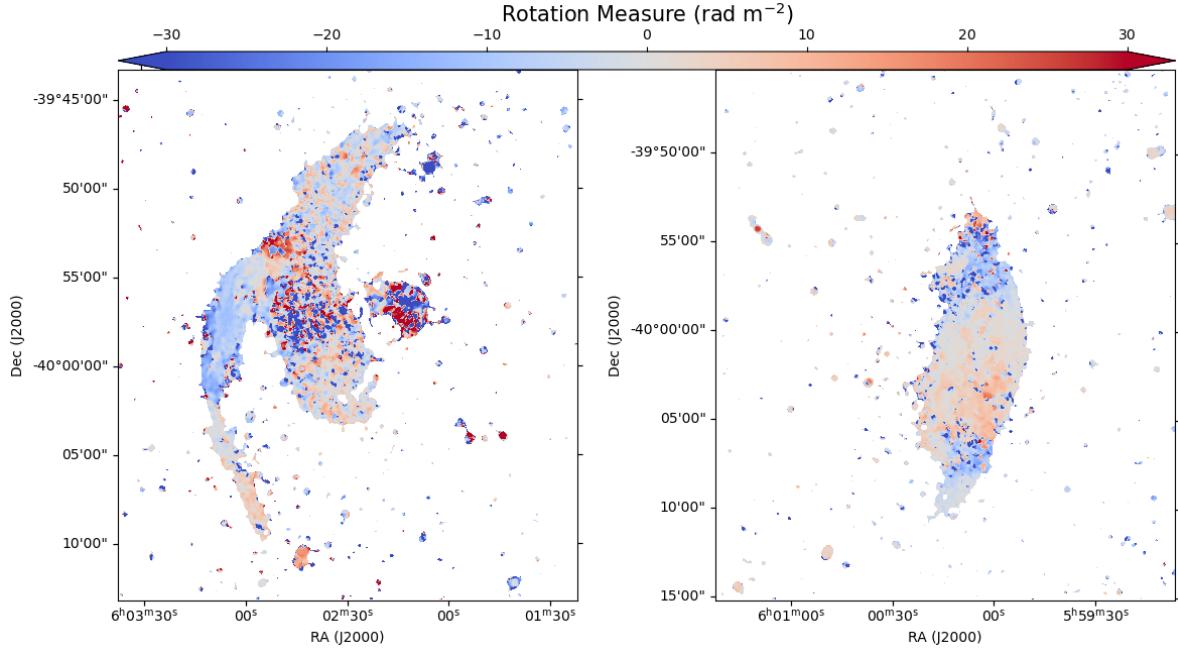
Fig. 13. Rotation measure map of Abell 3376 at a resolution of $15'' \times 15''$. For contrast, the colorbar is clipped at $\pm 30 \, \mathrm{rad \, m^{-2}}$. Local estimates for the Galactic Faraday rotation have been subtracted, with $\phi = 22 \, \mathrm{rad \, m^{-2}}$ for the eastern relic and $\phi = 32 \, \mathrm{rad \, m^{-2}}$ for the western relic. Only pixels with a total intensity above $3\sigma_I$ ($\sigma_I = 13 \, \mathrm{\mu Jy}$) are shown.

When comparing the predicted intrinsic polarization angles from Figs. 6 and 7, we observe greater deviations from the true values when multiple sources are included in each sample. This is expected due to the potential interference between signals that are close in ϕ. While this interference typically creates an unsolvable system of equations, the simulated dataset helps guide the model toward a reasonable solution.

In Fig. 8, we saw that depolarization effects, common in regular RM synthesis, can be reduced by learning the sensitivity window. However, caution should be taken when increasing the search range in ϕ, as this can produce false positives in the output. This occurs because the first part of Eq. (6) is RM sensitivity-dependent, as it includes the averaging operator.

6.2. Abell 3376

In this section we discuss the results from applying the model to observational data, including some implications and suggestions for future work.

While Fig. 10 captures the magnetic field orientation on large scales, as the overlaid projected magnetic field vectors are an average produced from a number of pixels, the small-scale features reveal a different pattern. For example, in Fig. 12, we see that while the magnetic field is ordered on small scales, it is also tangled and complex, indicating turbulent processes at play. In this figure we can more clearly see how the scales on which the magnetic field is ordered are reduced as we move inward into the relic. It thus becomes evident that the polarized fraction of the total flux should decrease due to beam depolarization effects. This effect is also seen in Fig. 14, where the regions far from the shock fronts show a decrease in polarized fraction. This decrease in polarized fraction aligns with expectations, as beam averaging over disordered magnetic fields at smaller scales leads to depolarization. Near the shock fronts, where the magnetic field is compressed and aligned, the polarized fraction remains high,

showing the shock's role in organizing the field. Further away, increased turbulence causes the reduction in polarization.

As the data are imaged using the CLEAN algorithm, the model image is convolved with a restoring beam to arrive at the final images. While the higher $8''$ resolution cubes could have revealed finer structures of the magnetic field, the $15''$ cubes were used in this work to improve the signal to noise at the regions of weak emission. Future telescopes such as the Square Kilometre Array (SKA) will allow us to study the magnetic field of radio relics in even greater detail, together with recovering lower surface-brightness features of polarized emission, due to greater sensitivity together with a higher angular resolution, thus reducing the amount of beam depolarization.

From the peaks in Fig. 9, we see that Gaussian peaks are favored by the model, most likely due to our Gaussian mixture model prior, rather than the true underlying signal. As physically realistic magnetic fields likely do not follow Gaussian profiles in ϕ (Bell et al. 2011), future studies should investigate the model's ability to recover non-Gaussian features. From the upper left plot in Fig. 9, we see that the flux increases toward higher $|\phi|$, which most likely is not the true reality of the relic magnetic field. These features might arise due to the model's confusion caused by the use of a restrictive prior. Including more realistic and non-Gaussian profiles in the simulated samples could thus make the model more robust and reduce the amount of unphysical features.

As mentioned in Sect. 4.1, due to the polarization map being contaminated by foreground emission, the threshold for training samples was set in total intensity. The result of this includes training on strongly depolarized signals, such as regions of turbulent magnetic fields, as well as unpolarized sources. While this can bias the model toward lower flux, and cause confusion due to the low S/N, it was preferred over including foreground signals in the training data, which would lead to poor reconstruction of the relic Faraday dispersion.
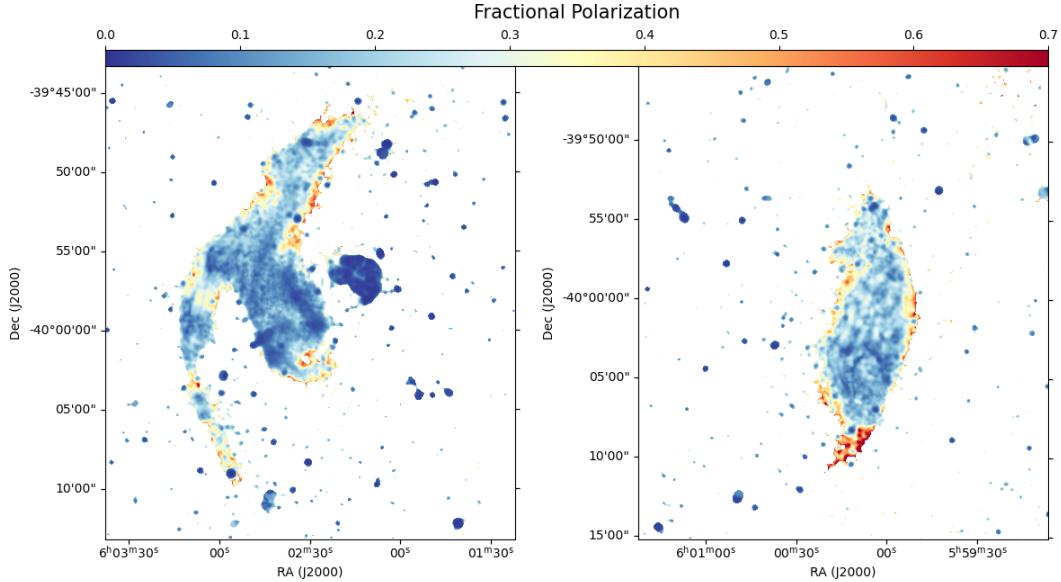
Fig. 14. Fractional polarization map of Abell 3376 at a resolution of $15'' \times 15''$. Only pixels with a total intensity above $3\sigma_I$ ($\sigma_I = 13\,\mu$Jy) are shown. The map has not been corrected for Rician bias.

While this work has focused on polarization data from the MGCLS, the model should perform similarly well in other frequency bands and with other telescopes. One only has to change the RM search range and resolution, as well as the FWHM of the sources that one expects to find. While it is possible to train the network on data from a range of instruments, we expect higher performance when tailored to a specific spectral band. Furthermore, the option of training on one dataset and inferring on another has not been studied in this work. This should in theory not be an issue, as long as the two datasets have a similar dynamic range and noise level. One could also create a training dataset from a variety of datasets which collectively represent a wide range of conditions, thereby improving the robustness and generalizability of the model.

6.3. Limitations and challenges

As discussed in Sect. 4.1, we do not add noise to the mock training data. A consequence of not including noise in the simulated samples is that the model sometimes produces false positives. Since the model is only trained at producing sparse representations, the model puts all the noise in peaks at locations that agree with the data. Examples can be seen in the upper middle and lower left plot in Fig. 9, where we have a low baseline, but small peaks separated from the main peaks. As the noise in the real and imaginary parts should ideally be Gaussian-distributed, the noise of their magnitude should follow a Rician distribution, that is, positive definite and asymmetric when the Gaussian noise have zero mean. Usually, before doing any quantitative analysis on the signal flux, one should therefore subtract the Rician bias from the Faraday spectrum. In this work, however, the opposite effect was observed, where the model systematically underestimated the total flux of signals. As this effect was not observed when noise was removed from the synthetic data, we believe there could be two reasons behind this. Either the model overestimates the noise level, leading to a conservative bias in the flux estimation, or some of the signal flux is put into the noise peaks.

The decision to omit noise was made for two main reasons: Firstly, accurately modeling the true noise statistics is difficult. Our model assumes that the noise is Gaussian with zero mean and uncorrelated in λ^2. However, in radio astronomical images, the noise is often correlated and non-Gaussian. Thus, adding purely Gaussian noise led to a decreased model performance, as the model struggled to transfer knowledge from simulations to real data and vice versa. We also attempted to create more realistic noise by extracting samples without significant signal from the data and then inserting a simulated signal into those samples. While this approach helped to reduce false positives in the outputs, it falls short of accurately reproducing the data.

The second reason concerns the loss function, which becomes problematic when incorporating noise in the simulations. The first component of Eq. (6) seeks to minimize the difference between the output in λ^2 and the observed data, including the noise. However, the second component aims to replicate the underlying signal, excluding the noise. Thus, the two loss terms compete against each other, and requires careful tuning of the β parameter in order for the model to benefit from both terms. By excluding noise in the simulated samples, both components of the loss function should aim to find a solution to the data, regardless of the noise statistics.

Furthermore, it is possible that Eq. (1) does not accurately capture the true measurement operation. One could thus draw inspiration from various techniques used in interferometric imaging, where terms are included to account for variations in the primary beam and the w-component. Similar effects in RM synthesis would include varying channel widths and bandwidth depolarization effects.

6.4. Summary

In this paper, we have developed a deep learning model to perform RM synthesis deconvolution. The model was trained using a Gaussian mixture model prior alongside observational data from MGCLS DR1. It is important to note that only 20% of the training data were simulated, meaning that while the simulated dataset guides the model to some extent, the majority of information was derived from the observational dataset, ensuring

the model remains primarily influenced by real-world data. We tested the model on simulations as well as observational data from the MeerKAT galaxy cluster data. The model was able to recover extended and high-RM components better than the traditional CLEAN algorithm. For future studies, it would be beneficial to compare the deep learning model with an algorithm designed to deconvolve Faraday thick components, such as a multi-scale variant of RMCLEAN or a multicomponent adaptation of *QU* fitting. When applied to data from the Abell 3376 galaxy cluster, the model was able to deconvolve Faraday spectra from various parts of the intracluster medium, and map the magnetic field structure of the cluster relics in fine detail. We find that while the dominant filament largely determines the magnetic field orientation, there are regions where this assumption does not hold. One should thus be cautious when using methods such as the synchrotron intensity gradient mapping, as studying clusters in full polarization can yield more comprehensive information. The model was also applied to five other datasets from MGCLS, out of which only two had previously been studied with MeerKAT in full Stokes. The results generally agree with those from previous studies, with the main difference being the great sensitivity of the MeerKAT telescope, providing a great amount of detail to, for example, the Abell 85 cluster where we have mapped the magnetic field structure of a complex phoenix. Furthermore, the model was able recover the magnetic field orientation from regions of weak emission from, for example, the Abell 3667 cluster, revealing ordered magnetic fields on megaparsec scales. The high accuracy and low computational cost of the model makes it a good candidate for forthcoming surveys, where an automated pipeline will be crucial to handle the large data sizes.

References

Alger, M. J., Livingston, J. D., McClure-Griffiths, N. M., et al. 2021, PASA, 38, e022
Andrecut, M., Stil, J. M., & Taylor, A. R. 2012, AJ, 143, 33
Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, ApJ, 935, 167
Bell, M. R., & Enßlin, T. A. 2012, A&A, 540, A80
Bell, M. R., Junklewitz, H., & Enßlin, T. A. 2011, A&A, 535, A85
Brentjens, M. A., & de Bruyn, A. G. 2005, A&A, 441, 1217
Brown, S., Bergerud, B., Costa, A., et al. 2019, MNRAS, 483, 964
Burn, B. J. 1966, MNRAS, 133, 67
Chibueze, J. O., Sakemi, H., Ohmura, T., et al. 2021, Nature, 593, 47
Chibueze, J. O., Akamatsu, H., Parekh, V., et al. 2023, PASJ, 75, S97
Cornwell, T. J. 2008, IEEE J. Sel. Top. Signal Process., 2, 793
de Gasperin, F., Rudnick, L., Finoguenov, A., et al. 2022, A&A, 659, A146
Domínguez-Fernández, P., Brüggen, M., Vazza, F., et al. 2021, MNRAS, 507, 2714
Duchesne, S. W., Johnston-Hollitt, M., & Bartalucci, I. 2021a, PASA, 38, e053
Duchesne, S. W., Johnston-Hollitt, M., Offringa, A. R., et al. 2021b, PASA, 38, e010
Dwarakanath, K. S., Parekh, V., Kale, R., & George, L. T. 2018, MNRAS, 477, 957
Enßlin, T. A., & Brüggen, M. 2002, MNRAS, 331, 1011
Enßlin, T. A., Biermann, P. L., Klein, U., & Kohle, S. 1998, A&A, 332, 395
Farnsworth, D., Rudnick, L., & Brown, S. 2011, AJ, 141, 191
Frick, P., Sokoloff, D., Stepanov, R., & Beck, R. 2010, MNRAS, 401, L24
George, L. T., Dwarakanath, K. S., Johnston-Hollitt, M., et al. 2015, MNRAS, 451, 4207
Giovannini, G., & Feretti, L. 2000, New Astron., 5, 335
He, K., Zhang, X., Ren, S., & Sun, J. 2015, ArXiv e-prints [arXiv:1512.03385]
Heald, G., Braun, R., & Edmonds, R. 2009, A&A, 503, 409
Hu, Y., Stuardi, C., Lazarian, A., et al. 2024, Nat. Commun., 15, 1006
Hutschenreuter, S., Anderson, C. S., Betti, S., et al. 2022, A&A, 657, A43
Ideguchi, S., Takahashi, K., Akahori, T., Kumazaki, K., & Ryu, D. 2014, PASJ, 66, 5
Johnston, S., Bailes, M., Bartel, N., et al. 2007, PASA, 24, 174
Jonas, J., & MeerKAT Team 2016, MeerKAT Science: On the Pathway to the SKA, 1
Kale, R., Dwarakanath, K. S., Bagchi, J., & Paul, S. 2012, MNRAS, 426, 1204
Kingma, D. P., & Ba, J. 2017, ArXiv e-prints [arXiv:1412.6980]
Knowles, K., Cotton, W. D., Rudnick, L., et al. 2022, A&A, 657, A56
Kumazaki, K., Akahori, T., Ideguchi, S., Kurayama, T., & Takahashi, K. 2014, PASJ, 66, 61
Li, F., Brown, S., Cornwell, T. J., & de Hoog, F. 2011, A&A, 531, A126
Miyashita, Y., Ideguchi, S., & Takahashi, K. 2016, PASJ, 68, 44
Oktay, O., Schlemper, J., Le Folgoc, L., et al. 2018, ArXiv e-prints [arXiv:1804.03999]
O'Sullivan, S. P., Brown, S., Robishaw, T., et al. 2012, MNRAS, 421, 3300
Pratley, L., & Johnston-Hollitt, M. 2020, ApJ, 894, 38
Pratley, L., Johnston-Hollitt, M., & Gaensler, B. M. 2021, PASA, 38, e060
Ronneberger, O., Fischer, P., & Brox, T. 2015, ArXiv e-prints [arXiv:1505.04597]
Rudnick, L., & Cotton, W. D. 2023, MNRAS, 522, 1464
Rudnick, L., Brüggen, M., Brunetti, G., et al. 2022, ApJ, 935, 168
Schmidt, K., Geyer, F., Fröse, S., et al. 2022, A&A, 664, A134
Slee, O. B., & Reynolds, J. E. 1984, PASA, 5, 516
Slee, O. B., Roy, A. L., Murgia, M., Andernach, H., & Ehle, M. 2001, AJ, 122, 1172
Sun, X. H., Rudnick, L., Akahori, T., et al. 2015, AJ, 149, 60
Van Eck, C. L. 2018, Galaxies, 6, 112
van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2
Wang, R., Chen, Z., Luo, Q., & Wang, F. 2023, ArXiv e-prints [arXiv:2305.09121]

Appendix A: Results from other galaxy clusters observed with MeerKAT

We also applied our deep learning deconvolution model to other clusters from the MGCLS DR1. For some of these clusters RM data have already been published and for others we show the polarization data for the first time. In this appendix we present results from some interesting sources. A more detailed physical interpretation is going to be presented in separate papers.

Abell 3186   The cluster Abell 3186 (also called MCXC J0352.4-7401) has been studied by, for example, Duchesne et al. (2021a) using the Murchison Widefield Array (MWA) and ASKAP telescopes. Most recently it has been observed by Hu et al. (2024) with MeerKAT, using the SIG method to study the magnetic field structure of the relics. In Fig. A.1 the two radio relics in Abell 3186 are shown. We see that in both relics the magnetic field aligns well with the orientation of the relics, particularly near the shock front. Overall, the results align well with Hu et al. (2024), as the polarization vectors generally follow the orientation of the dominant filament. However, at the inner edge of the southwestern relic, some differences are noticeable. In Fig. A.1, the magnetic field is oriented toward the inner region of the cluster, while the SIG method consistently results in field vectors perpendicular to the local gradient. A similar pattern is observed in the north-eastern relic, although it is less pronounced. These findings suggest that SIG might not capture certain details revealed by polarization studies, particularly when the magnetic field exhibits unexpected orientations. Furthermore, studying the polarization properties at a pixel-by-pixel basis allow us to obtain a unbiased result, independent on the surrounding region.

Abell 168   Next we show the results from the cluster Abell 168. This double relic cluster has previously been studied by, for example, Dwarakanath et al. (2018) using the Karl G. Jansky Very Large Array (VLA) and the Giant Meterwave Radio Telescope (GMRT). However, polarization studies have not been conducted of the cluster. In Fig. A.2 the cluster is shown together with the projected magnetic field orientation. While an arc-shaped relic is located, the second relic found by Dwarakanath et al. (2018) is not visible. The reason for this is that regions far from the pointing center are blanked in the MGCLS enhanced products, which is the case for the relics in this cluster. In the relic detected here it is striking how uniform the magnetic field orientation is throughout the relic. At least at our angular resolution it does not show the variance that is observed in some of the other relics.

Abell 3667   The cluster Abell 3667 has recently been studied extensively in full polarization with MeerKAT L-band by de Gasperin et al. (2022). The authors discovered that the relics are composed of a network of synchrotron filaments with varying spectral and polarization properties. These are likely linked to multiple regions of particle acceleration and localized magnetic field enhancements. In Fig. A.3 we show the magnetic field orientation of the two radio relics. We see that the magnetic field is ordered on scales comparable to the relic extension, and that the main contributor to the field orientation is the local orientation of the dominant filament. When comparing our results with de Gasperin et al. (2022), we find that the magnetic field in our analysis appears to remain ordered even in regions of weak emis-sion, whereas de Gasperin et al. (2022) indicates turbulent magnetic fields in these areas. Since both studies are based on the same set of visibilities, the differences might stem from either the CLEANing process to a higher resolution or the thresholding methods used by de Gasperin et al. (2022). Furthermore, the authors use a Faraday depth sampling of $2\,\mathrm{rad\,m^{-2}}$, as opposed to $1\,\mathrm{rad\,m^{-2}}$ in this work. As the regions of weak emission are expected to be Faraday depolarized, a small change in ϕ would potentially result in a great difference in the observed polarization angle. Additionally, the authors used a Högbom algorithm similar to RMCLEAN to produce the polarization maps, but it is uncertain if this is the cause of the divergent results.

Abell 85   In Fig. A.4 we show a cutout from the cluster Abell 85, displaying some interesting sources. It is worth mentioning that we did not include the radio halo at the center of the cluster identified by Knowles et al. (2022), as its emission is very weak. Most notably in our map to the north-west we find a complex phoenix (revived fossil plasma source), which has previously been studied by, for example, Slee & Reynolds (1984), Giovannini & Feretti (2000) and Duchesne et al. (2021b). However, the classification varies between a phoenix and a relic. It is interesting to note that in the northern and southern filaments the magnetic field vectors follow the geometry of the filaments. Furthermore, southeast of the phoenix, we identify an AGN with wide angle jets, whose magnetic field orientation align well with the outflow direction. Slee et al. (2001) could not determine the field vectors because they could not Faraday-derotate their data. Also they only saw the central and northern parts of this phoenix. But they found significant variation in the polarization fraction in A85, ranging from 35% in the northern part to 10% in the central part of the phoenix (they label it as southeastern arc). In Fig. A.5 the fractional polarization map from the MGCLS is shown. The map is produced from the mean of $\sqrt{Q^2 + U^2}$ over the entire band. As the resulting polarized flux is not corrected for Rician bias, the polarized fractions should be considered as upper limits. Similar to Slee et al. (2001) we find a polarization fraction of about 10% in the central part of the phoenix, but with some regions below 5%. In the filaments the polarized fraction is generally higher, around 10-15%, with some hotspots with a fractional polarization between 20% and 30%.

In Fig. A.6 the magnetic field orientation of the inner region of the phoenix is shown, in order to highlight the vortex-like structures. We see that the magnetic field of this region is very complex and in some regions ordered on just the scale of the angular resolution. In the regions of the bright torus structures, the field lines appear to follow the orientation of the filament, while where there is weak emission the magnetic field is more chaotic. These rapid rotations of the magnetic field are likely the cause of the depolarization effects seen in the central region of Fig. A.5. Enßlin & Brüggen (2002) proposed that the torus structures seen in Fig. A.6 are formed by light radio plasma moving through a shock wave. At the shock front, the ram pressure of the pre-shock gas and the thermal pressure of the post-shock gas are in balance. When the lighter radio plasma comes in contact with the shock front, the ram pressure is reduced at the point of contact. This causes the post-shock gas to expand into the volume occupied by the low-pressure radio plasma. This process disrupts the radio plasma, eventually forming a torus structure, similar to a smoke-ring. The similarity between the simulations of Enßlin & Brüggen (2002), and the magnetic field in Fig. A.6 strengthen the argument that these toroidal structures are a result of such a compression scenario.
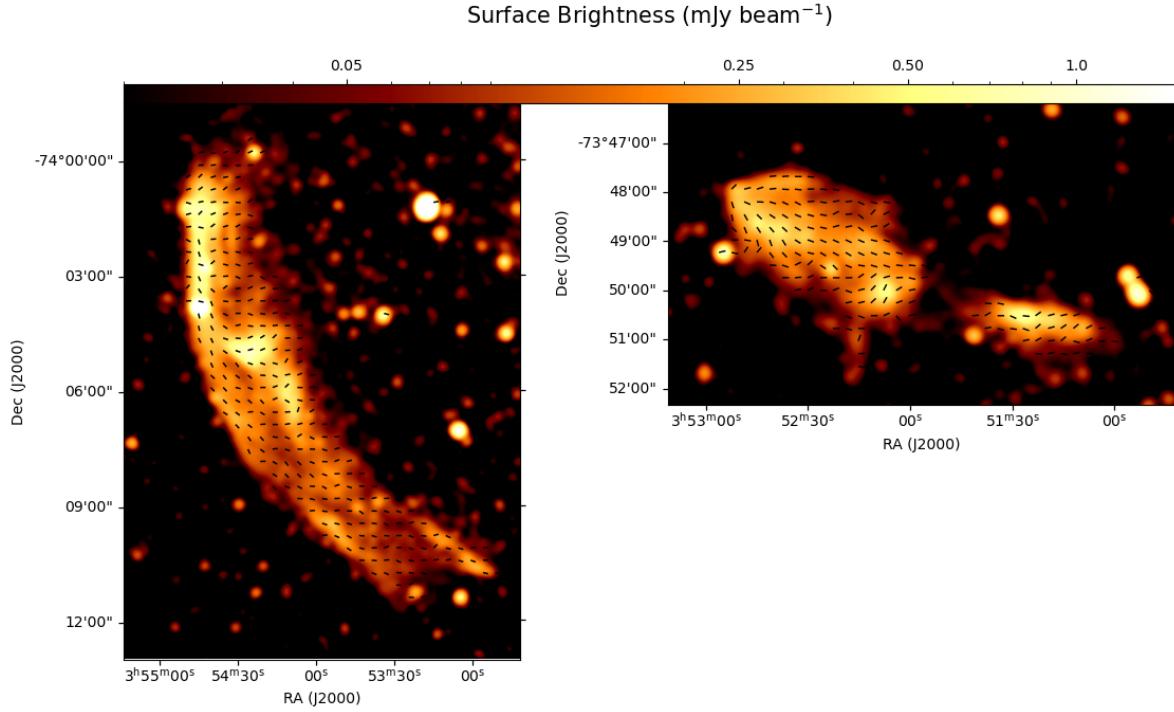
Fig. A.1. Radio emission from the Abell 3186 galaxy cluster relics at a resolution of $15'' \times 15''$, overlaid with the polarization *B*-vectors.
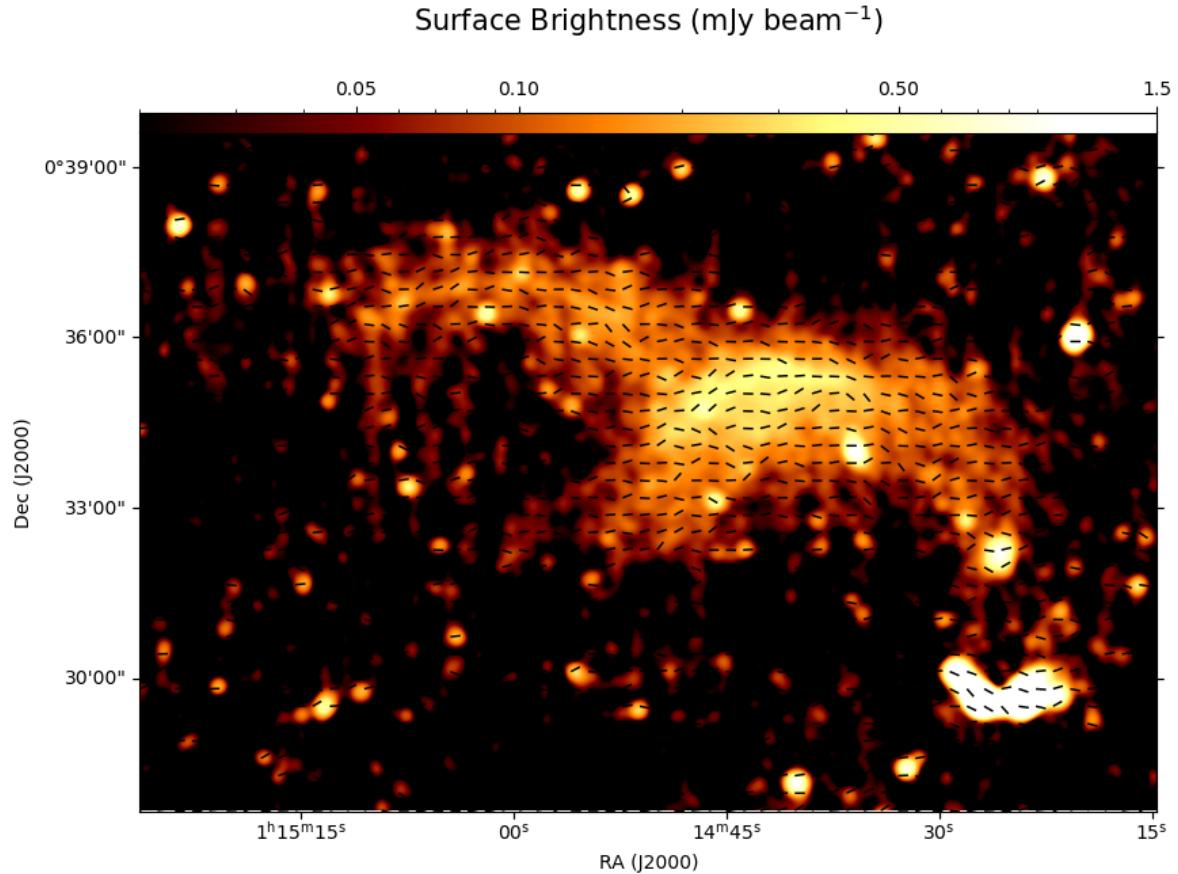

Fig. A.2. Radio emission from the Abell 168 galaxy cluster relic at a resolution of $15'' \times 15''$, overlaid with the polarization *B*-vectors.
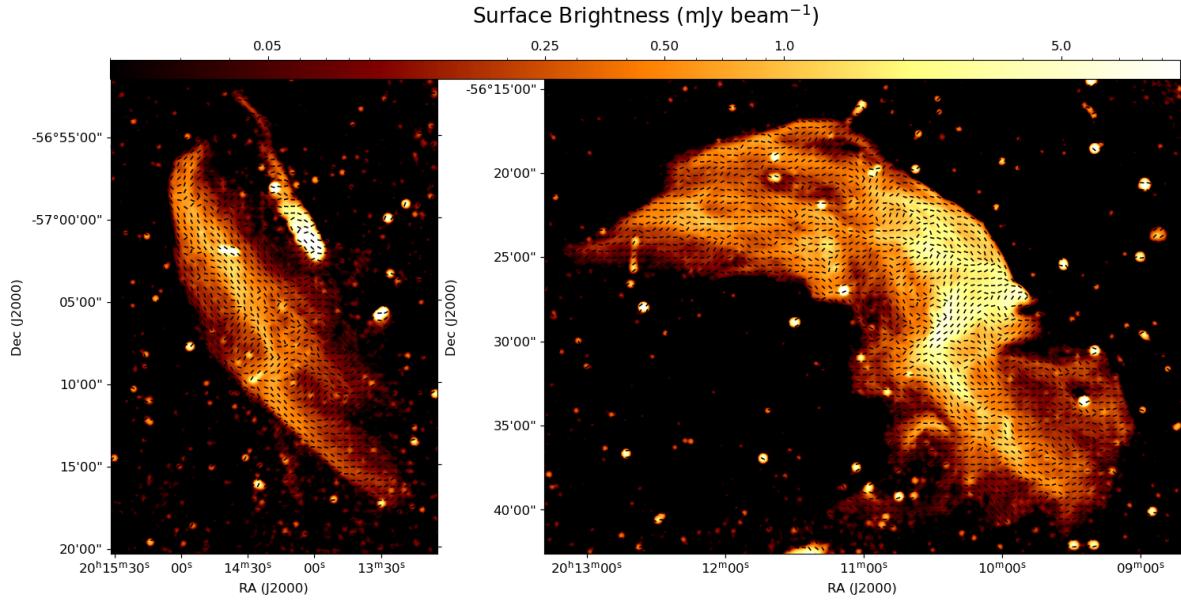
Fig. A.3. Radio emission from the Abell 3667 galaxy cluster relics at a resolution of $15'' \times 15''$, overlaid with the polarization B-vectors.
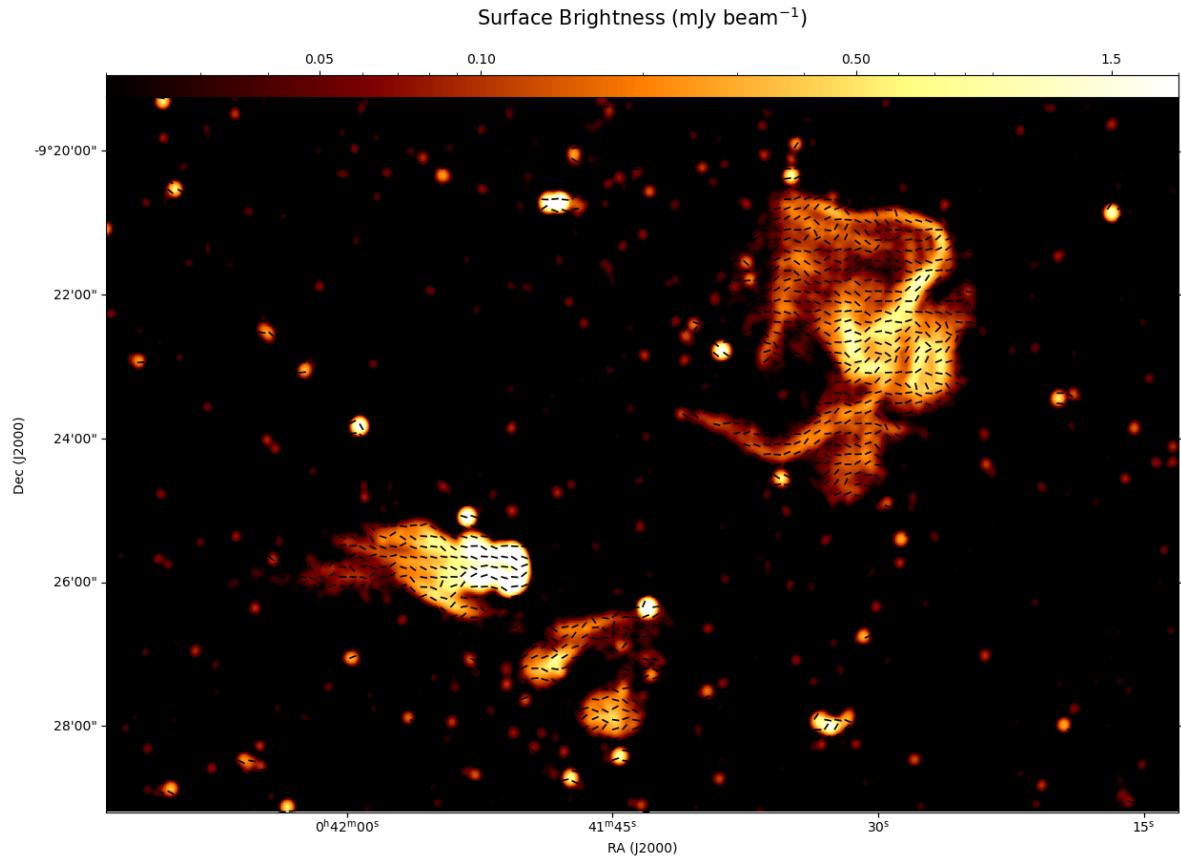
Fig. A.4. Radio emission from the galaxy cluster Abell 85 at a resolution of $8'' \times 8''$, overlaid with the polarization B-vectors.

Abell 194 Finally, in Fig. A.7 we show the radio emission together with the polarization B-vectors for two radio galaxies in the galaxy cluster Abell 194. The same dataset has previously been studied by Rudnick et al. (2022), including an extensive RM study of the two radio galaxies, together with a map of the magnetic field of the long eastern filament. The eastern filaments are around 50% polarized, with no detectable net RM with

respect to the Galactic foreground, and only small rms variations ($9\,\mathrm{rad\,m^{-2}}$) along their length. This enables a mapping between Faraday depth and distance along the line-of-sight. The magnetic field vector is shown to follow the length of the filament, supporting the interpretation that these are magnetic flux tubes. In Fig. A.8 a comparison between this work and Rudnick et al. (2022) is shown for the filament. We see that our results are
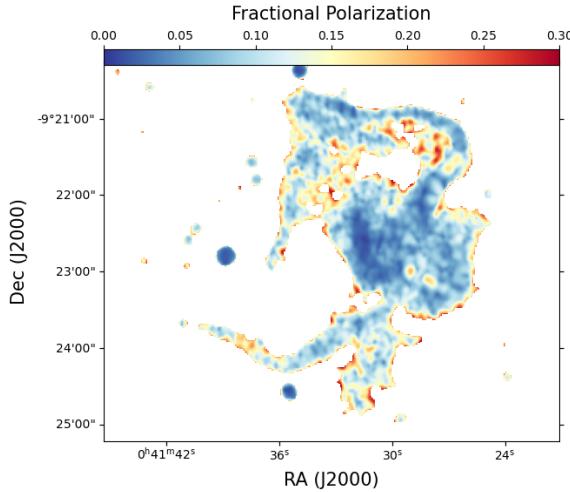
Fig. A.5. Fractional polarization map (908-1656 MHz) of the phoenix in Abell 85, at a resolution of $8'' \times 8''$. Only pixels with a total intensity above $5\sigma_I$ ($\sigma_I = 12\,\mu$Jy) are shown.
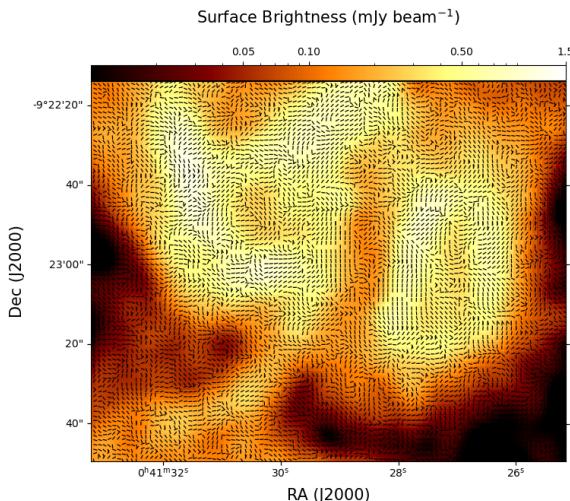

Fig. A.6. Radio emission from the central region of the phoenix in Abell 85 at a resolution of $8'' \times 8''$, overlaid with the polarization B-vectors.

similar to those of Rudnick et al. (2022). The field vectors in our work seem to vary more smoothly over the filament. However, this might just be to different vector lengths and number of samples used to produce such vectors.
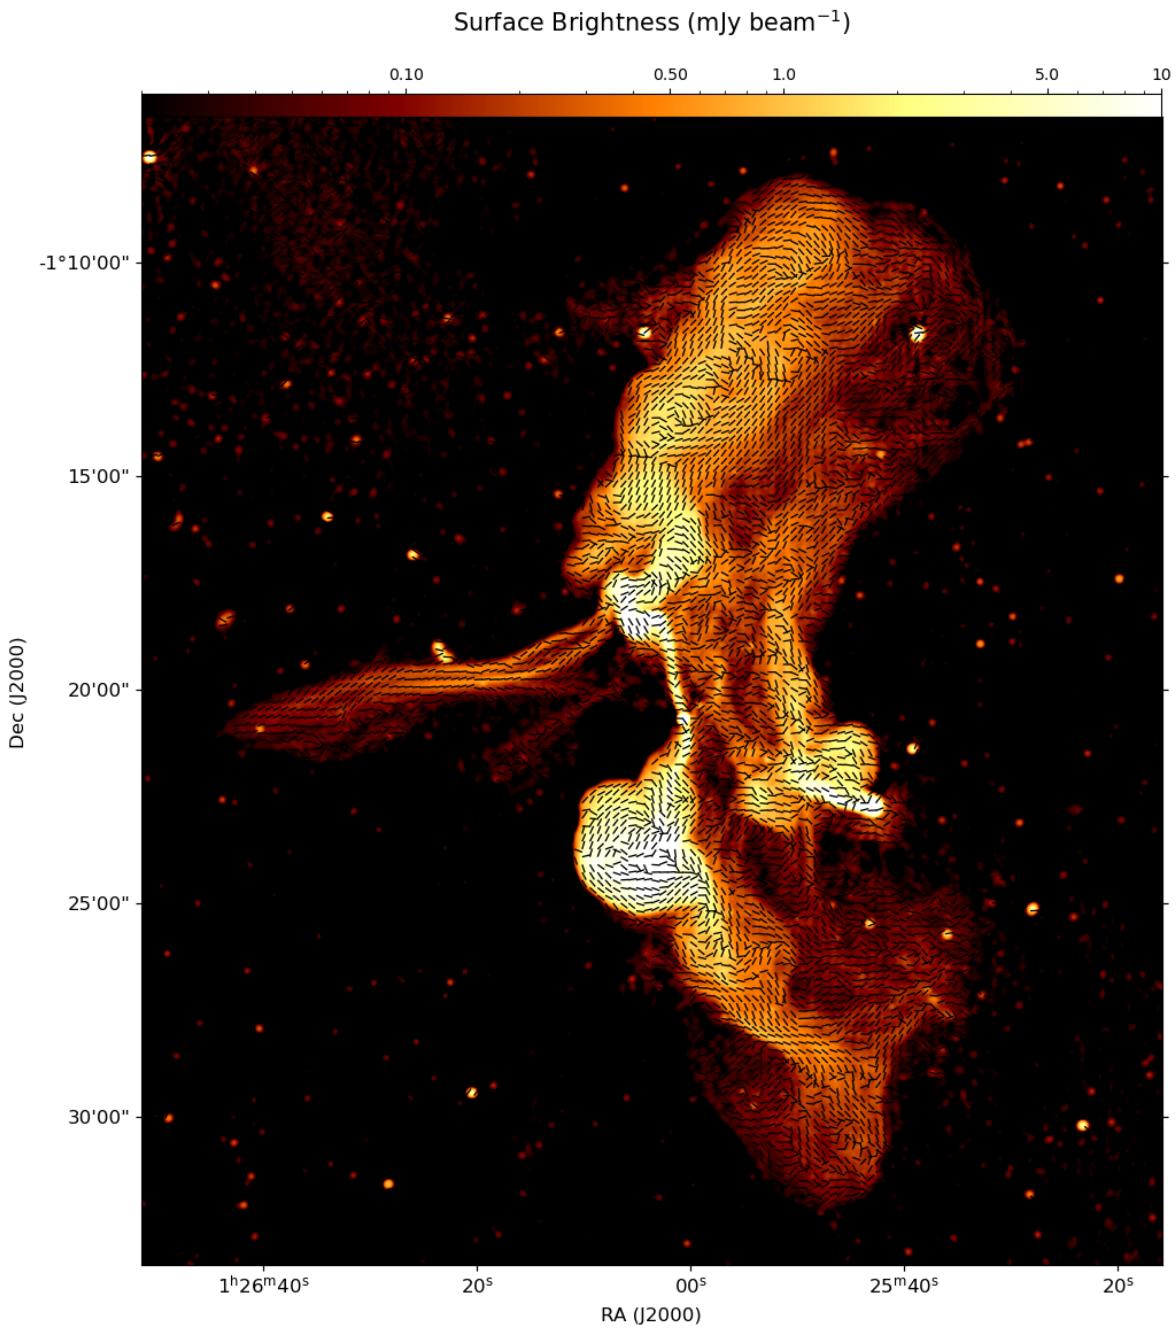
Fig. A.7. Radio emission from two radio galaxies in the galaxy cluster Abell 194 at a resolution of 8″×8″, overlaid with the polarization *B*-vectors.
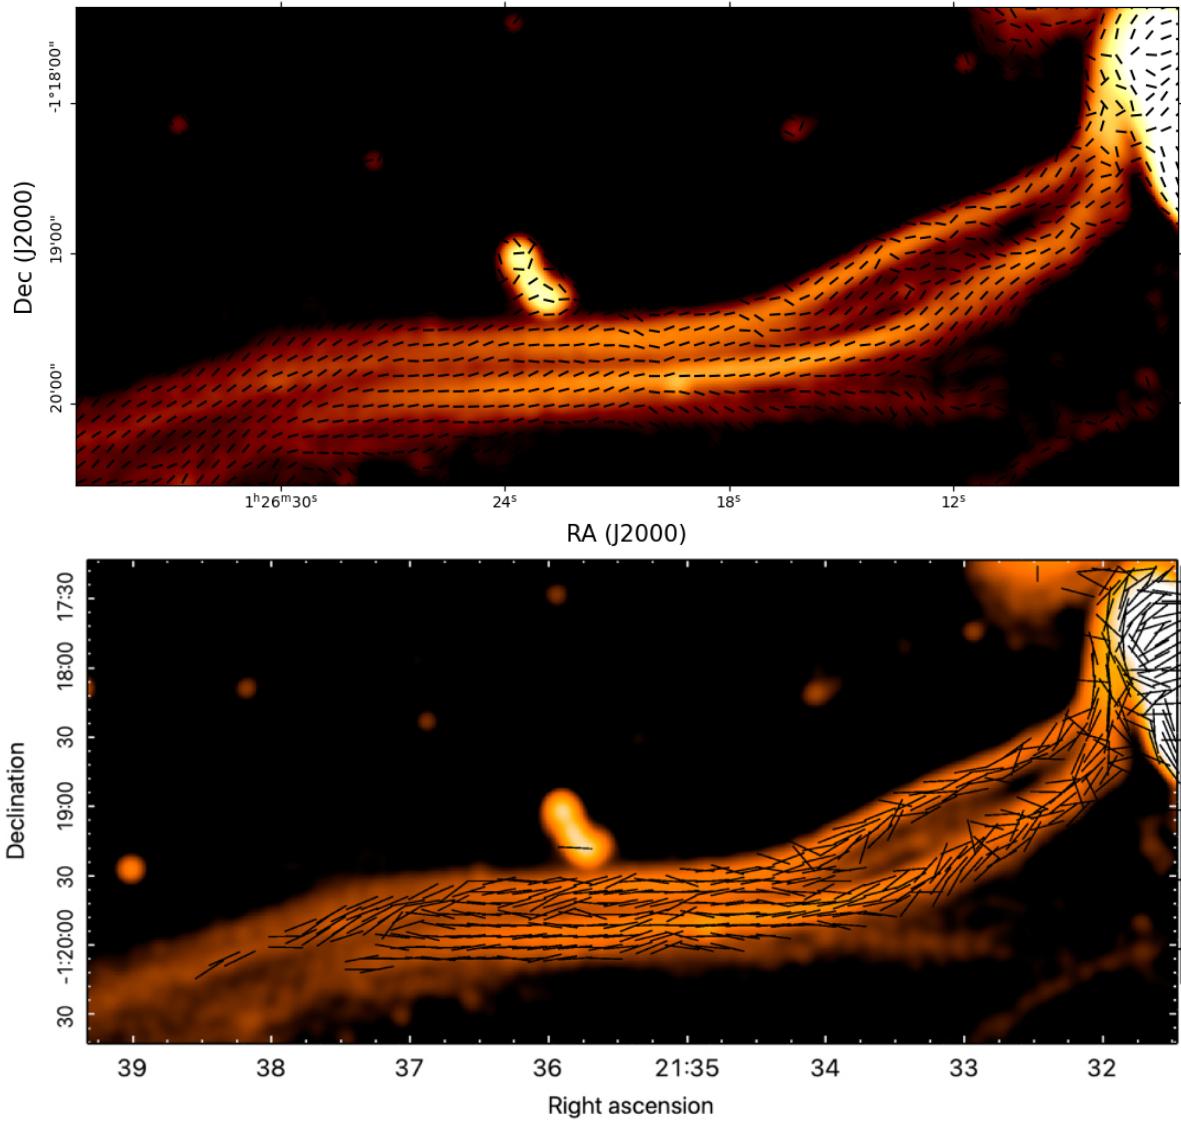
Fig. A.8. Comparison between this work (top) and Rudnick et al. (2022) (bottom), of the magnetic field orientation in the long eastern filament.