

Mapping diffuse radio sources using TUNA: a transformer-based deep learning approach

N. Sanvitale,¹★ C. Gheller,¹ F. Vazza^{1,2}, A. Bonafede,² V. Cuciti,² E. De Rubeis^{1,2}, F. Govoni,³ M. Murgia^{1,3} and V. Vacca^{1,3}

¹*Istituto di Radio Astronomia, INAF, Via P. Gobetti 101, I-40129 Bologna, Italy*

²*Dipartimento di Fisica e Astronomia, Università di Bologna, Via P. Gobetti 92/3, I-40129 Bologna, Italy*

³*Osservatorio Astronomico di Cagliari, INAF, Via della Scienza 5, I-09047 Cagliari, Italy*

Accepted 2025 July 10. Received 2025 June 14; in original form 2025 April 21

ABSTRACT

Vision transformers are used via a customized TransUNet architecture, which is a hybrid model combining transformers into a U-Net backbone, to achieve precise, automated, and fast segmentation of radio astronomy data affected by calibration and imaging artefacts, addressing the identification of faint, diffuse radio sources. Trained on mock radio observations from numerical simulations, the network is applied to the LOFAR Two-meter Sky Survey data. It is then evaluated on key use cases, specifically megahaloes and bridges between galaxy clusters, to assess its performance in targeting sources at different resolutions and at the sensitivity limits of the telescope. The network is capable of detecting low surface brightness radio emission without manual source subtraction or re-imaging. The results demonstrate its groundbreaking capability to identify sources that typically require reprocessing at resolutions 4–6 times lower than that of the input image, accurately capturing their morphology and ensuring detection completeness. This approach represents a significant advancement in accelerating discovery within the large data sets generated by next-generation radio telescopes.

Key words: methods: data analysis – techniques: image processing – galaxies: clusters: intracluster medium – large-scale structure of Universe.

1 INTRODUCTION

Radio observations of galaxy clusters provide invaluable insights into the properties of the non-thermal components of the intracluster medium, revealing the presence of cosmic-ray electrons and magnetic fields across volumes of several cubic megaparsecs. These components may extend beyond the cluster’s virial radius, potentially permeating the cosmic web (see e.g. Ryu et al. 2003; Pfrommer et al. 2006; Ryu et al. 2008; Vazza et al. 2025). Radio emission, primarily found in massive, dynamically active galaxy clusters, is traditionally classified as giant radio haloes, mini haloes, and radio relics based on their morphology, origin, and cluster location (for reviews Feretti et al. 2012; Brunetti & Jones 2014; Van Weeren et al. 2019; Wittor 2023). Modern low frequency radio-interferometric observations, characterized by unprecedented sensitivity, revealed synchrotron-emitting bridges connecting pairs of massive clusters in a pre-merger phase (see e.g. Govoni et al. 2019; Botteon et al. 2020; Pignataro et al. 2024), possibly formed through turbulence in the intracluster region (Brunetti & Vazza 2020), providing the first evidence of cosmic structures that trace the large-scale structure of the Universe. These observations have also allowed the discovery of a new class of large diffuse radio sources, outside the typical radius of ‘classical’ radio haloes ($\geq 0.4R_{500}$ out to R_{500} , where

R_{500} is the spherical radius enclosing a total matter overdensity 500 times greater than the cosmic mean density at the corresponding redshift), referred to as ‘megahaloes’ (Cuciti et al. 2022). Bridges and megahaloes are predicted to be rare compared to classical radio haloes and relics, since they might originate during specific stages of the cluster evolution and in presence of particular merger configurations (Brunetti & Vazza 2020). The detection of the few occurrences present in current and forthcoming radio surveys is thus of paramount importance.

The common feature of these observations is that they target extended sources, challenging to detect given their low surface brightness emission combined with their large angular scale. Characterizing their morphology and properties requires overcoming the limitations imposed by the constrained observational sensitivity due to factors like radio frequency interference, ionospheric effects, and limited dynamic range. Furthermore, even in an ideal scenario of perfectly calibrated data, detecting extended emission necessitates accurate subtraction of compact sources that dominate the interferometric signal. These sources, if not removed properly, can leave sidelobes or residuals interfering with the detection of low-surface-brightness structures. This typically requires re-imaging the data at lower spatial resolutions, which helps enhance sensitivity to large angular scales. However, this re-imaging process is highly computationally demanding and time consuming and it must be carefully tuned, as it can lead to the suppression of real features, to the amplification of random and confusion noise.

* E-mail: nicoletta.sanvitale@inaf.it

The methodological challenges are worsened by the volume of the data involved. Such observations result from the processing of huge data sets generated by state-of-the-art radio telescopes such as the Low-Frequency Array (LOFAR) (van Haarlem et al. 2013), the Karoo Array Telescope (MeerKAT) (Jonas & Team 2016), the Murchison Widefield Array (MWA) (Tingay et al. 2013), and the Australian Square Kilometre Array Pathfinder (ASKAP) (Hotan et al. 2021), whose unprecedented sensitivity, resolution, frequency coverage, and field of view significantly increase data volume. They are precursor to the hundreds of petabytes of data that the Square Kilometre Array Observatory (SKAO¹) is expected to generate annually. Efficient numerical solutions that are capable of performing tasks such as radio source detection, segmentation, and flux calculation in a fast, accurate, and fully automated manner are necessary to manage such large amounts of complex data, minimizing the necessity for human supervision, which becomes impractical for large data sets, and eliminating the need for reprocessing or re-imaging the data through traditional pipelines. In this work, we explore the adoption of a customized version of the TransUnet, hereafter referred as *TransUnet Network for radio Astronomy (TUNA)* to radio observations. The network is trained using mock radio images generated from numerical simulations, then applied to various observational targets selected from the second data release of the LOFAR Two-metre Sky Survey (LoTSS-DR2) and the Planck Sunyaev–Zel’dovich (PSZ2) cluster catalogue (Botteon et al. 2022). This catalogue serves as a reference for the statistical analysis of TUNA’s capabilities and performance. We also compare these results with those obtained using a method based on the original U-Net architecture (Stuardi et al. 2024), in order to explore the differences and similarities between the two approaches. The TUNA network is subsequently evaluated on several additional cases, specifically megahaloes and bridges (Govoni et al. 2019; Botteon et al. 2020; Cuciti et al. 2022), to assess its performance in targeting sources at the limits of LOFAR sensitivity and in complex radio environments.

2 RELATED WORK

Machine and deep learning (ML/DL) have the potential to effectively address the requirements of large-scale radio astronomy data processing, as demonstrated by the number of applications in different scientific domains published over the last decade (see e.g. Radovic et al. 2018; Vamathevan et al. 2019; Cavanagh, Bekki & Groves 2021; Rolnick et al. 2022). In particular, the authors have explored the use of DL in radio astronomy, focusing on convolutional neural networks for automated detection of sources (Gheller, Vazza & Bonafede 2018), Denoising Autoencoders for image quality enhancement (Gheller & Vazza 2022), and fully convolutional U-Net architecture (hereafter R-UNet) for semantic segmentation of diffuse radio emission (Stuardi et al. 2024).

The advent of the transformer architecture (Vaswani et al. 2017) has had a profound impact on various domains of AI, particularly on natural language processing (NLP), representing the foundation for breakthrough Large Language Models, such as BERT (Devlin et al. 2019), GPT (Radford et al. 2018), Gemini (Gemini Team 2023) and DeepSeek (DeepSeek-AI et al. 2024). At their core, transformers use a self-attention mechanism (Vaswani et al. 2017), which allows the model to focus on different parts of the input to capture key relationships and extract intrinsic features from data.

¹<https://www.skao.int/>

Transformers quickly extended to other fields, first computer vision. This paper investigates the use of vision transformers (ViT) (Dosovitskiy et al. 2020) to enhance the effectiveness of machine learning methods for processing radio interferometric data. ViTs process images as patches, using self-attention to capture relationships across them. TransUnet (Chen et al. 2021), introduced in 2021 for medical image segmentation, is a pioneering hybrid model integrating ViTs into the U-Net architecture (Ronneberger, Fischer & Brox 2015), combining the extensive contextual awareness of the former with the localized feature extraction abilities of the latter. This enables the identification of extended targets with large-scale correlations, while accurately capturing small-scale features, making it an ideal solution for the identification and segmentation of diffuse radio sources.

3 METHODS

This section presents an overview of the TUNA network architecture, the training data employed, and its performance on mock observations.

3.1 The TUNA architecture

TUNA relies on the TransUnet network (Chen et al. 2021). This is a hybrid architecture which integrates the vision transformer module, as introduced by Dosovitskiy (2020), into the U-Net architecture (Ronneberger et al. 2015) to address the constraints of the limited receptive field of convolutional operations, which hinder U-Net’s ability to capture long-range contextual dependencies.

Fig. 1 shows that TransUnet retains the encoder–decoder structure of U-Net but employing a hybrid CNN-Transformer encoder. In this architecture, CNN, specifically ResNet50 (He et al. 2016), serves as a feature extractor, producing feature maps that encapsulate fine-grained spatial information. If the input image has dimensions $H \times W \times C$ (height, width, and number of channels), the output feature map dimensions are given by:

$$H' = \frac{H}{P}; \quad W' = \frac{W}{P}, \quad (1)$$

where P is the downsampling factor of the CNN backbone ($P = 16$ for ResNet50).

Each feature map is then divided into patches of size $(P_s \times P_s)$, which are linearly projected into D -dimensional embeddings. The TransUnet architecture assumes $P_s = 1$. To encode spatial information, position embeddings are learned and added:

$$\mathbf{z}_0 = [\mathbf{x}_p^1 \mathbf{E}; \mathbf{x}_p^2 \mathbf{E}; \dots; \mathbf{x}_p^N \mathbf{E}] + \mathbf{E}_{\text{pos}}, \quad (2)$$

where $\{\mathbf{x}_p^i \in \mathbb{R}^{P_s^2 \cdot C'} | i = 1, \dots, N\}$ are the flattened 2D patches, $N = \frac{HW}{P_s^2}$ is the number of image patches (i.e. the input sequence length), C' is the number of feature maps, $\mathbf{E} \in \mathbb{R}^{(P_s^2 \cdot C') \times D}$ is the patch embedding projection, and $\mathbf{E}_{\text{pos}} \in \mathbb{R}^{N \times D}$ denotes the position embedding.

The resulting sequence of embedding vectors serves as input to the transformer encoder. The transformer encoder (Vaswani et al. 2017) consists of a stack of $n = 12$ identical layers. Each layer sequentially processes the input, with the output of one layer serving as the input for the next. Each layer has two blocks. The first is a multi-head self-attention mechanism (MSA), and the second is a simple, position-wise fully connected feed-forward network (MLP). There are residual connections around both the MSA and the MLP blocks. Layer normalization (LN) is applied after each residual connection

3.2 Training data

TUNA was trained on synthetic observations of diffuse radio emission, produced from cosmological simulations. Images have been produced to closely resemble actual LOFAR observations, following the methodology outlined in Gheller & Vazza (2022) and Stuardi et al. (2024). Here we summarize the main steps of the procedure.

The cosmological simulations were performed using ideal magnetohydrodynamics (MHD) with the Eulerian code *ENZO* (Bryan et al. 2014). The simulation box has a uniform spatial cell resolution of 41.65 kpc (comoving) and it encompasses a comoving volume of 100^3 Mpc^3 . A uniform primordial magnetic seed field of $B_0 = 0.1 \text{ nG}$ (comoving) was initialized everywhere in the box at the start of the simulation ($z = 50$) and was later evolved under ideal MHD conditions until $z = 0$. The simulation assumed a standard Λ CDM cosmological model, with density parameters $\Omega_b = 0.0478$, $\Omega_{\text{DM}} = 0.2602$, $\Omega_\Lambda = 0.692$, and a Hubble parameter of $H_0 = 67.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

In post-processing, we calculated the emission at 150 MHz from relativistic electrons accelerated by cosmic shocks at different redshifts, assuming that only shocks can accelerate relativistic particles through diffusive shock acceleration and generate synchrotron radio emission (e.g. Brown 2011). The model of synchrotron emission proposed by Hoeft & Brüggén (2007) is employed in this context,

while we did not account for additional radio emission generated by the reacceleration by turbulence on relativistic electrons, which likely has a key role in the formation of radio haloes, bridges, or megahaloes (Brunetti 2011; Brunetti & Vazza 2020; Nishiwaki et al. 2024). However, the morphology and emissivity of diffuse radio sources can loosely resemble even those of radio haloes, despite the different particle acceleration mechanisms likely at work in radio haloes (e.g. Van Weeren et al. 2019; Lee et al. 2023).

The creation of the training set of synthetic images for TUNA, based on the 3D emission model from cosmological simulations, followed two main basic steps. First, we integrated the emission along the line of sight within the same volume at four different redshift snapshots (roughly equally spaced from $z = 0.02$ to $z = 0.15$).

Next, every single synthetic sky model was generated by progressively stacking maps referred to an increasing redshift. In order to fully cover the $\approx 640 \text{ Mpc}$ distance out to $z = 0.15$, we extracted maps multiple times from the same simulated redshifts, applying an appropriate cosmological corrections factor to simulate an increasing redshift of the sources, similar to Vazza et al. (2021). As done in Gheller & Vazza (2022), we limited our analysis of simulated light cones up to $z \approx 0.15$ since we do not expect a drastic change in the radio emission including larger redshifts (see also Hodgson et al. 2021).

Table 1. Performance metrics for TUNA and R-UNet calculated for the positive class (diffuse emission). The second and third columns present scores for the simulated data at 6 arcsec and the last two columns for the simulated data at 20 arcsec.

| Metric | TUNA 6 arcsec | R-UNet 6 arcsec | TUNA 20 arcsec | R-UNet 20 arcsec |
|-----------|-----------------|-----------------|-----------------|------------------|
| Recall | 0.50 ± 0.01 | 0.39 ± 0.01 | 0.62 ± 0.02 | 0.53 ± 0.01 |
| Precision | 0.61 ± 0.02 | 0.64 ± 0.01 | 0.76 ± 0.02 | 0.73 ± 0.01 |
| IoU | 0.38 ± 0.01 | 0.30 ± 0.01 | 0.52 ± 0.01 | 0.43 ± 0.01 |
| Accuracy | 0.97 ± 0.01 | 0.96 ± 0.01 | 0.97 ± 0.01 | 0.97 ± 0.01 |




Figure 3. The two rows show examples of Clean images at 6 arcsec used as input for the network (left panels), sky maps with surface brightness exceeding $10^{-8} \text{ Jy pixel}^{-1}$ (our ground truth, right panels), predictions from TUNA (mid-left panels), and predictions from R-UNet (mid-right panels), all shown at the 0.5 confidence level.

By applying random rotations to the various redshift slices, over 500 independent light cones were generated. The images have a resolution of 2000×2000 pixels, covering a field of view of $1.1^\circ \times 1.1^\circ$ each, with a nominal angular resolution (pixel size) of 2 arcsec. These images are referred to as *sky images*.

To generate mock observations, the synthetic visibilities corresponding to a LOFAR HBA observation of 8 h have been calculated using the ‘predict’ mode of the WSCLEAN with the sky images as model images (Offringa et al. 2014; Offringa & Smirnov 2017). At this point, a random Gaussian noise with rms similar to LotSS noise is added to the visibilities using the noise.py script in the LOSITO software package.² Finally, imaging is performed using the WSCLEAN software, adopting Uniform weighting and correcting for the primary beam. The use of these parameters resulted in a restoring beam of $5.9 \text{ arcsec} \times 5.1 \text{ arcsec}$. Deconvolution is carried out using the clean method and the automasking option with a 5σ threshold. The resulting images are referred to as *clean images*. Along with the corresponding sky images, they compose the training and test data sets for TUNA. We will refer to this images as the 6 arcsec training data set. From the sky images, the ground truth is generated by assigning a value of 0 to all pixels with surface brightness less than $10^{-8} \text{ Jy pixel}^{-1}$ (Class 0) and a value of 1 to all other pixels (Class 1), creating the binary masks that have to be reproduced by the network. The resulting data set is unbalanced, with Class 1 pixels accounting for only 3 per cent of the total pixel count on average. The same procedure has been repeated to create a 20 arcsec training data set, convolving the sky images with a Gaussian function, to achieve a 20 arcsec resolution.

3.3 Network training, optimization, and performance analysis

The following methodology is used to train and assess TUNA on mock observations. One hundred 2000×2000 pixel images are divided into $N_T \times N_T$ pixel tiles (generated adding suitable padding) and used for training the network. Ten per cent of the data are used as a validation set to assess the network’s performance during the training process. Forty additional images are used as a testing data set to evaluate the final performance of the network. The training has been performed using both 6 and 20 arcsec data sets.

To determine the optimal network configuration, we performed multiple trainings using various hyperparameter settings. More specifically, we have set the learning rate $\mu = [0.01, 0.005, 0.001, 0.0001]$, batch size $B = [12, 24, 48]$, and tile size $N_T = [224, 512]$. The learning rate controls the step size at which a model updates its weights during training; the batch size defines the number of training samples processed before updating the model’s weights in each iteration; the tile size defines the dimension of a subregion into which the input image is divided for training. The network performance has been evaluated on a pixel basis using four metrics: Recall, Precision, Intersection over Union (IoU), and Accuracy, calculated by comparing the inferred mask to the ground truth. The four metrics are defined as follows, computed for the positive class (diffuse emission):

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (8)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (9)$$

$$\begin{aligned} \text{IoU} &= \frac{\text{Intersection}}{\text{Union}} \\ &= \frac{\sum (\text{Predicted Mask} \cap \text{Ground Truth Mask})}{\sum \text{Predicted Mask} \cup \text{Ground Truth Mask}} \end{aligned} \quad (10)$$

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}, \quad (11)$$

where TP (True Positives) and FP (False Positives) represent the number of pixels correctly and incorrectly classified as diffuse emission, while TN (True Negatives) and FN (False Negatives) correspond to pixels correctly and incorrectly classified as non-emitting, respectively. Recall measures the fraction of true sources correctly identified, while Precision quantifies the proportion of predicted sources that are TP. IoU evaluates the overlap between predicted and ground truth masks, and Accuracy represents the overall fraction of correctly classified instances across both diffuse emission and non-emitting regions.

The performance of the various configurations is summarized in Appendix A. The best configuration was chosen based on the values of the reported metrics and by visual inspection of the results on the test images. For both 6 and 20 arcsec data sets, it is achieved with the following setup: $\mu = 0.005$, $B = 24$, $T = 512$.

Adopting such ideal setup, the uncertainty associated with the training process was estimated by recalculating the performance metrics through five repeated training runs, following a slight variation of the K-fold cross-validation approach (Kohavi 1995). In each training session, a different subset comprising four-fifths of the total training set is used. The resulting models are then applied for inference on a fixed test set consisting of the same forty images.

The validation loss, that is, the error calculated in the validation data set after each training epoch (E), is presented in Fig. 2 for the




Figure 4. Surface brightness distribution for the simulated data set at 6 arcmin, considering pixels above $10^{-8} \text{ Jy pixel}^{-1}$, accounting for the pixels identified by TUNA and R-UNet.

6 arcsec data set. The figure displays the cross entropy (CE) and dice curves, along with the total loss, which is the average of the two losses. The three curves are presented for Class 0 combined with Class 1, as well as for Class 1 only.

The combined loss curves indicate that training convergence is reached within the first 50 epochs. However, the Class 1 losses keep decreasing up to 200 epochs, indicating that additional training is

²<https://github.com/darafferty/losito/tree/master>

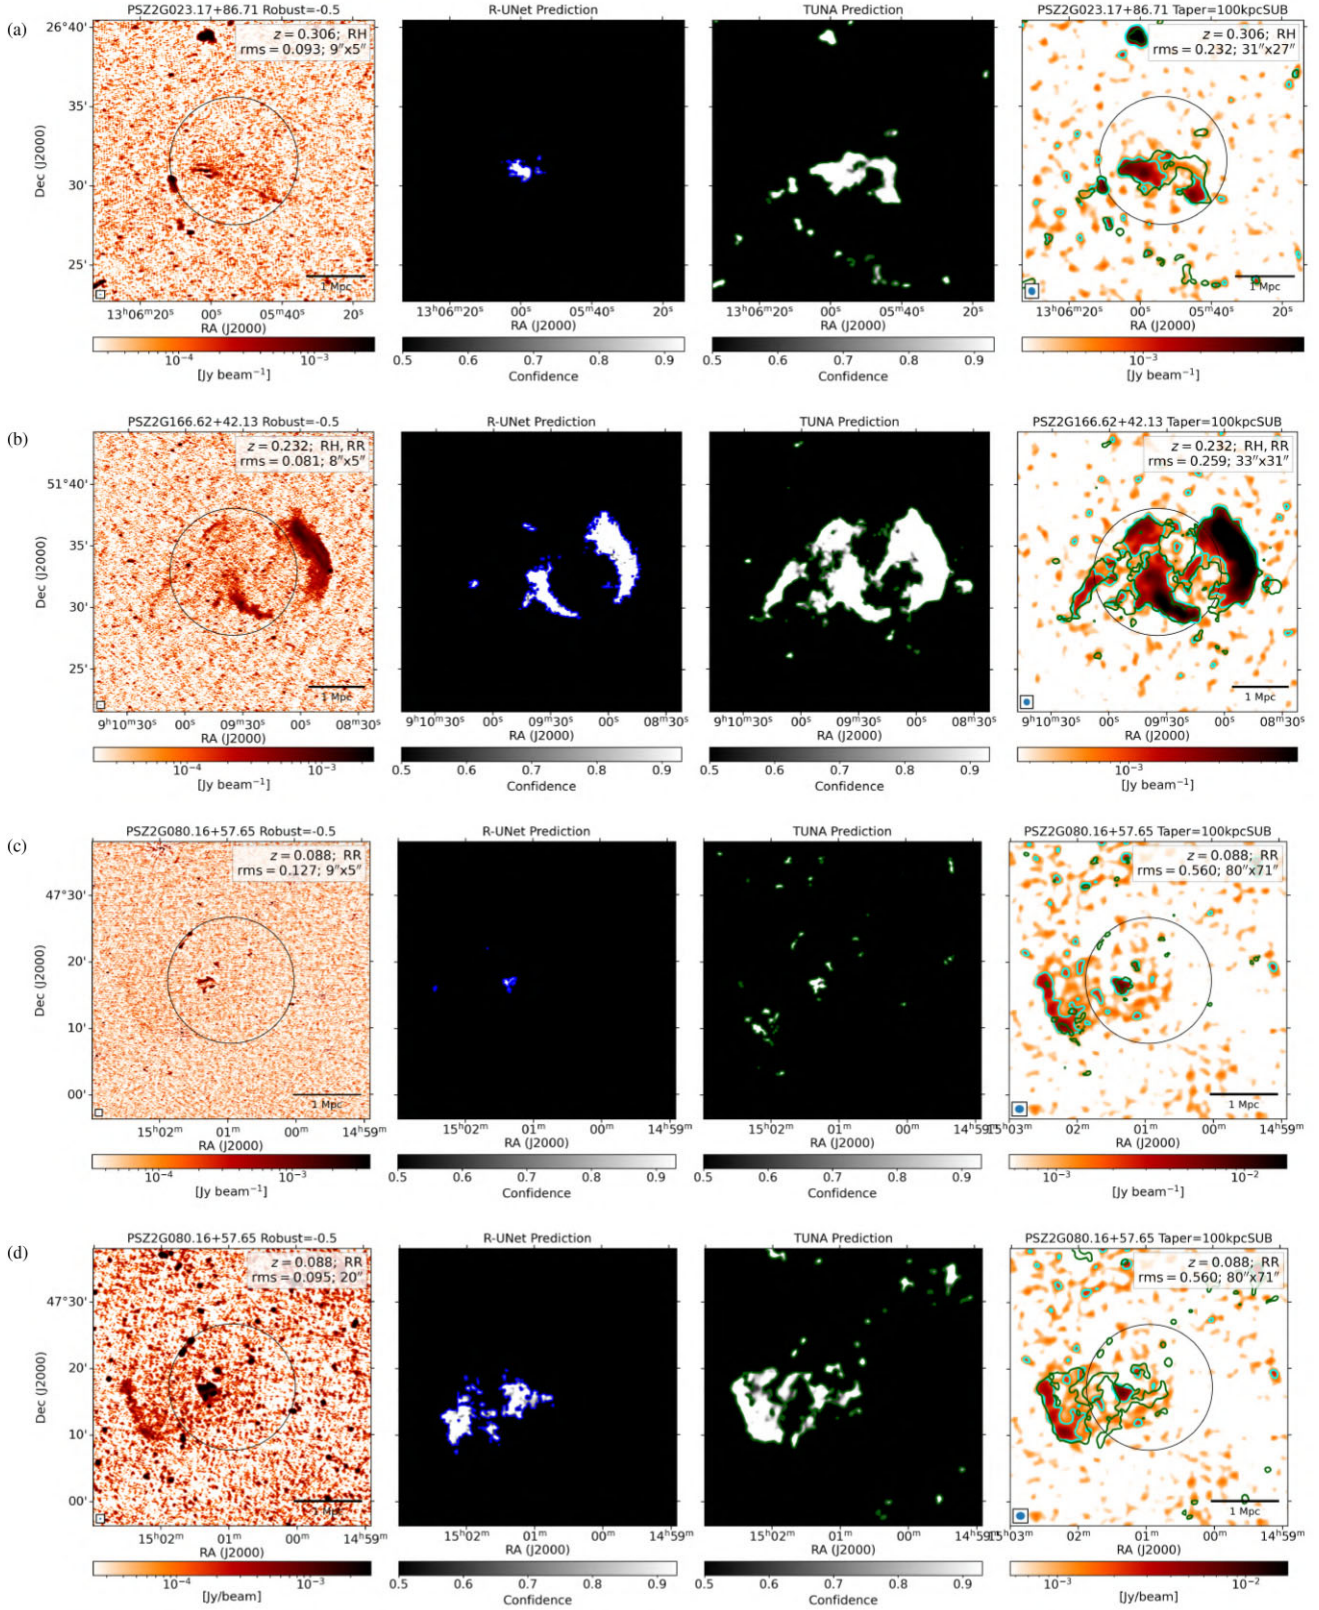


Figure 5. Left panel: radio images used as input for the networks. Central panels: prediction from the networks at a confidence level ≥ 0.5 with blue and green contours representing the confidence level = 0.5 for the R-UNet and TUNA networks, respectively. Right panel: 100 kpc tapered source-subtracted image with 3σ level cyan contour and TUNA prediction green contour superimposed. The colour scale in the left panel is logarithmically stretched from 0.25 to 30σ ; in the right panel, from 1 to 30σ . The redshift (z), image noise (rms, in mJy beam^{-1} units), classification, and image resolution of each cluster are reported in the top-right corner. The circle denotes the R_{500} cluster radius.

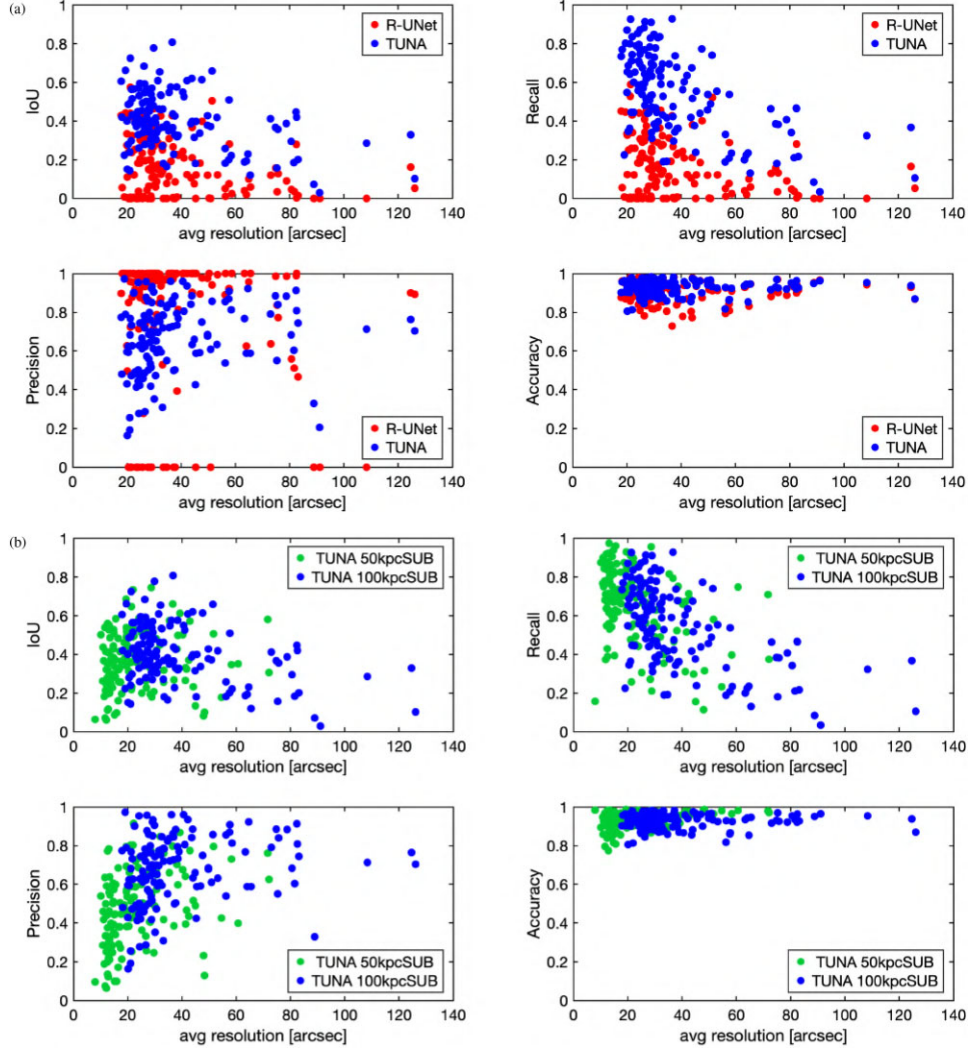


Figure 6. Panel (a): performance scores of TUNA and R-UNet as a function of the average resolution of the 100 kpc source-subtracted tapered images. Panel (b): performance metrics of TUNA as a function of the average resolution of the reference images, including 100 and 50 kpc source-subtracted tapered images. The average resolution is calculated as $(B_{\text{maj}} \times B_{\text{min}})^{1/2}$ where B_{maj} and B_{min} represent the major and minor axes of the beam, respectively.

Table 2. Performance metrics for TUNA and R-UNet calculated for the positive class (diffuse emission). The second and third columns present scores for the LoTSS-DR2/PSZ2 data set.

| Metric | TUNA | R-UNet |
|-----------|-----------------|-----------------|
| Recall | 0.61 ± 0.18 | 0.20 ± 0.16 |
| Precision | 0.65 ± 0.19 | 0.77 ± 0.35 |
| IoU | 0.43 ± 0.13 | 0.19 ± 0.15 |
| Accuracy | 0.93 ± 0.04 | 0.92 ± 0.05 |

needed for the network to effectively learn to classify Class 1 pixels. This occurs because the data set is highly imbalanced with Class 1 constituting approximately 3 percent of all pixels. To address this issue, we applied class weighting in the loss function, assigning a higher weight to the underrepresented Class 1, according to the following equation:

$$\mathcal{L} = \mathcal{L}_{\text{Class0}} + w_1 \mathcal{L}_{\text{Class1}}, \quad (12)$$

where w_1 is the weighting factor (the values [1, 2, 10] have been tested) and $\mathcal{L}_{\text{ClassN}}$ is the contribution to the loss of Class N. However, this strategy proved ineffective in improving the model’s ability to classify Class 1 pixels (see Appendix A). The only effective solution was to extend the training to 200 epochs. The same approach also proved effective for the 20 arcsec training set.

Table 1 presents the values of the metrics obtained for Class 1 using the optimal set-up, alongside the corresponding values achieved with the R-UNet model (Stuardi et al. 2024). For the 6 arcsec model TUNA has a Recall of 0.50 ± 0.01 , whereas R-UNet has a Recall of 0.39 ± 0.01 , indicating a higher effectiveness of the former to identify emitting pixels. Furthermore, TUNA demonstrates superior performance in terms of IoU compared to R-UNet, suggesting that it is more effective also at accurately capturing the morphology of sources. The R-UNet demonstrates slightly higher Precision. This is primarily because R-UNet tends to identify the brightest pixels within the distribution, which are more reliably matched to the ground truth. As a result, its predictions align more accurately with the most




Figure 7. Surface brightness distribution for the LoTSS-DR2/PSZ2 observational catalogue considering pixels above 3σ (left), accounting for the pixels identified by TUNA and R-UNet.

prominent structures in the data set. Finally, the two networks exhibit comparable Accuracy, approaching unity. This reflects the fact that both can precisely identify Class 0 pixels, which predominate the input images.

At a lower resolution of 20 arcsec, all scores improve as a result of a larger restoring beam, which increases the sensitivity to diffuse emission enhancing the detectability of fainter regions, although at the expenses of the finer details along the source edges.

Fig. 3 presents two examples of mock images used as network input and ground truth for training, as well as predictions given by TUNA and R-UNet at the 0.5 confidence level. Both networks demonstrate the ability to segment the input image even in the presence major artefacts. As also observed above, TUNA is capable of identifying a greater number of pixels at low surface brightness and demonstrates an improved capacity to accurately reproduce the source’s actual morphology.

In Fig. 4, the distributions of the surface brightness of the pixels exceeding 10^{-8} Jy pixel $^{-1}$ along with those selected by the TUNA and by the R-UNet predictions are shown. For surface brightness above 10^{-5} Jy pixel $^{-1}$, both networks successfully detect all source pixels. At lower values, an increasing fraction of the emitting pixels cannot be detected in the clean image, due to noise and artefacts. However TUNA identifies a larger number of sources than R-UNet and detects a large fraction (more than 50 per cent) of the emitting pixels down to 10^{-7} Jy pixel $^{-1}$. Noticeable, at even lower surface brightness, the network can detect a significant fraction of the emitting pixels.

It is important to mention that all the training and testing were conducted with high performance computing resources on the Booster partition of the Leonardo supercomputer³ at the CINECA Italian Supercomputing centre. The Leonardo system comprises 3456, 32-cores, Intel Xeon Platinum 8358 CPU nodes equipped with 4 NVIDIA Tesla Ampere 100 GPUs. Each CPU node is equipped with 512 GB of DDR4 memory, while the GPU has 64 GB of HBM2 memory. The network can efficiently leverage GPU acceleration thanks to the support provided by the PyTorch framework and can exploit parallel computing for hyperparameters space exploration.

The computational performance on a single node of Leonardo is 2.9 sec per image for the inference and 88.5 sec per epoch for the training, using 2000×2000 pixel images.

4 RESULTS

Although trained on mock observations from cosmological simulations, TUNA proves to successfully generalize to real observational data, effectively distinguishing between artefacts and diffuse sources. In order to assess its performance on observational data, the network has been applied to multiple use cases, namely the Planck Catalogue from the LOFAR Two-meter Sky Survey (LoTSS), radio bridges in merging galaxy clusters and mega haloes. The observational data are converted to logarithmic scale and normalized between 10^{-8} and 10^{-2} .

4.1 Planck catalogue from LoTSS-DR2

The LoTSS-DR2 data is used to evaluate the performance of TUNA on real, observational data. LoTSS-DR2 covers the 27 percent of the northern sky observed in the range 120–168 MHz (the nominal central frequency is 144 MHz). Each LoTSS pointing is 8 h long, the typical resolution is ~ 6 arcsec, and the median root-mean-square (rms) noise is $\sigma \sim 0.083$ mJy beam $^{-1}$. The survey includes low-resolution (20 arcsec) images with median sensitivity $\sigma \sim 0.095$ mJy beam $^{-1}$ (Shimwell et al. 2022a).

We used as a benchmark the radio observations of the 309 clusters in PSZ2 that lie within the 5634 deg 2 covered by the LoTSS-DR2. This data set, hereafter referred to as LoTSS-DR2/PSZ2, represents the most comprehensive study to date of diffuse synchrotron emission in the intracluster medium. We employed data available from the LoTSS repository⁴ produced by Botteon et al. (2022).

The LoTSS-DR2/PSZ2 clusters have a redshift between $0.016 < z < 0.9$ and are classified as radio haloes (RH), relics (RR), uncertain sources (U), candidate radio halo, or candidate radio relic (cRH or cRR) in the absence of X-ray data. No diffuse emission (NDE) classifies objects that do not exhibit diffuse emission, though they may display lobes or tails from active galactic nucleus (AGN) in the field. Not applicable (N/A) refers to objects that cannot be properly classified due to inadequate data quality. Clusters classified as N/A and those lacking redshift were excluded from analysis, resulting in a data set of 246 objects. For each cluster, the images obtained using the Briggs weighting scheme (Briggs 1995) with $\text{Robust} = -0.5$ and no uv taper, with a typical resolution of ~ 6 arcsec, were used as input for TUNA and R-UNet.

The performance of the network was evaluated by comparing its predictions with source-subtracted images, in which sensitivity to extended emission was enhanced through the removal of discrete sources from uv data and subsequent re-imaging using a Gaussian uv taper. We used images generated with $\text{Robust} = -0.5$ and Gaussian uv taper corresponding to 100 kpc at the cluster redshift, ensuring the detection of diffuse radio emission while accounting for resolution differences across redshifts (see Appendix B). In this data set, radio relics have been observed within the maximum distance of $2.2R_{500}$ (Jones et al. 2023), hence $2 \times 2.2R_{500}$ cut-outs have been produced to perform our analysis.

Panels (a), (b), and (c) of Fig. 5 show examples of the results generated by the networks, along with the input and reference images. TUNA accurately reproduces the morphology of the elongated radio

³<https://www.hpc.cineca.it/systems/hardware/leonardo/>

⁴https://lofar-surveys.org/planck_dr2.html




Figure 8. On the left, the A399-A401 ridge; on the right, the Abell 1758 bridge. Top panels: radio images used as input for TUNA. Central panels: TUNA predictions at 0.5 confidence level; cyan contours show the 0.9 confidence level for A399-A401 and 0.5 for A1758. Bottom panels: low resolution reference images with 3σ contours (white) and confidence contours (cyan) superimposed. The image resolution is shown on the top of each panel.

halo in the cluster PSZ2G023.17+86.71, as shown in the 100 kpc source-subtracted tapered image. R-UNet only detects the brightest areas of the radio halo. The PSZ2G166.62+42.13 cluster hosts a double radio relic, one prominent in the north-west and the other symmetrically positioned. Two additional fainter and smaller relics in the East and north-east directions have been identified. Additionally, a radio halo with low surface brightness is observed (Rajpurohit et al. 2024). R-UNet detects the two main relics and the brightest region of the radio halo, whereas TUNA reproduces precisely all the complex structures of the system. The PSZ2G080.16+57.65 is a low surface brightness giant radio relic-candidate radio halo system (Cavanagh et al. 2021). Both networks detect the radio halo, but only TUNA

can partially identify the relic. Panel (d) shows that when using the low-resolution (20 arcsec) image (Shimwell et al. 2022a) as input, both networks successfully detect the diffuse emission associated with the radio halo and relic. However, TUNA identifies a larger region of diffuse emission and better traces the shape of the relic. This larger emission detected by TUNA is supported by observations in Cavanagh et al. (2021) using a 90 arcsec source-subtracted tapered image.

The masks generated by TUNA effectively replicate the size and morphology of the radio sources, identifying diffuse emission that becomes detectable only at much lower resolution than the input data and after subtraction of the embedded sources. This is confirmed by




Figure 9. The LoTSS-DR2 pointing P128+37 enclosing the cluster Abell 697 (one of the four megahaloes). The top row shows the high resolution (6 arcsec) radio observation (left panel) and the inferred TUNA mask (right panel). The bottom row shows the low resolution (20 arcsec) radio observation (left panel) and the inferred TUNA mask (right panel).

examining the entire image set from the catalogue LoTSS-DR2/PSZ2 (available on-line, see the Data Availability Section). The network can remove the artefacts produced by the imaging process and exhibit the ability to detect a large number of radio galaxies within the field.

The effectiveness of segmentation has been assessed using the four metrics defined in Section 3.3: Recall, Precision, Intersection over Union (IoU), and Accuracy, calculated by comparing the inferred mask to the ground truth. Due to the unavailability of the true surface brightness distribution on the sky, the analysis used a mask derived from pixels with surface brightness above 3σ in subtracted tapered images as a proxy for ground truth. This approach is considered a conservative approximation aimed at minimizing residual noise. We considered the pixels inside a circular area of radius R_c , centred on the cluster’s centre, where $R_c = 2.2R_{500}$ for relics and $R_c = 1.1R_{500}$ for all other scenarios. The results obtained for each cluster are presented in Appendix B.

Panel (a) of Fig. 6 shows the relationship between the evaluation scores with the average resolution of the 100 kpc source-subtracted tapered images (131 clusters, having neglected NDEs). The IoU and Recall metrics indicate that TUNA outperforms R-UNet in effectively capturing source morphology and ensuring detection completeness. Although R-UNet exhibits higher precision than TUNA, this is attributed to sub-segmentation, where only smaller portions of sources are identified, leading to fragmented detections that fail to capture the full structure of the sources. The accuracy, significantly affected by the large fraction of non-emitting regions, is close to unity.

Panel (b) of Fig. 6 displays the evaluation metrics of TUNA, using both the 100 and 50 kpc source-subtracted tapered images from

Botteon et al. (2022) as reference for the ground truth masks. Two separate taperings allows the assessment of TUNA’s performance at different resolutions, pinpointing the specific resolution at which the network’s segmentation mask is comparable with diffuse emission in lower resolution images. The results indicate that TUNA achieves the best match within the resolution range of 20 to 40 arcsec. In this range, IoU is higher, and Recall and Precision exhibit their best balance, with precision dropping at higher resolution and recall diminishing at lower resolution. The decreased precision is due to the smaller size of the sources with respect to the corresponding masks. This demonstrates TUNA’s ability to identify faint diffuse emissions at a resolution 4–6 times lower than that of the input data (~ 6 arcsec).

The average evaluation metrics within the optimal range of 20–40 arcsec (see Table 2) confirm that TUNA outperforms R-UNet, with IoU more than twice as high and Recall more than three times greater. R-UNet has slightly higher Precision and both networks exhibit consistent Accuracy.

Fig. 7, show the surface brightness for pixels above 3σ , alongside the surface brightness identified by TUNA and R-UNet predictions that overlap this threshold mask. Above 10^{-2} Jy beam $^{-1}$, both networks detect the majority of source pixels. TUNA continues to identify a significant fraction (over 50 per cent) of the signal above 3σ extending down to 10^{-3} Jy beam $^{-1}$, the missing detections primarily due to the presence in the images of noise and artefacts exceeding 3σ . Below 10^{-3} Jy beam $^{-1}$, TUNA keeps recognizing a significant fraction of emitting pixels.

The capability of TUNA in identifying diffuse emission at resolutions 4–6 times lower than input data was also evaluated using




Figure 10. The four megahaloes ZwCl 0634.1+4750, Abell 665, Abell 697, and Abell 2218. Left column: LOFAR 144 MHz radio images used as input for TUNA. Right column: low resolution, source subtracted images with 3σ contours (white) and TUNA confidence contours (cyan) superimposed. The confidence level is at 0.9 for ZwCl 0634.1+4750 and 0.5 for the remaining cases. The image resolution is shown on the top of each panel. The colour scale in the left panel uses a logarithmic stretch from 0.5 to 200σ , while the right panel uses 0.5 to 50σ .

low-resolution 20 arcsec images, as shown in panel (d) of Fig. 5. The analysis compared masks generated using 20 arcsec images with those from 100 kpc source-subtracted tapered images with resolution ≥ 70 arcsec (15 clusters in the entire data set, available

as on-line material, see Section 5). While the sample size limits statistical reliability, the results suggest that TUNA can detect diffuse emission observable at resolutions between 60 and 120 arcsec (see also Section 4.3).

4.2 The A399-A401 Ridge and the Abell 1758 Bridge

Several recent low-frequency observations have shown the presence of diffuse radio emission between interacting cluster pairs in a pre-merger phase (Govoni et al. 2019; Botteon et al. 2020).

Govoni et al. (2019) observed a ~ 3 Mpc ridge of radio emission connecting the pre-merging system made by Abell 399 and Abell 401 (A399-A401). The ridge was observed with the LOFAR high band antenna (HBA) at a central frequency 140 MHz. We use the 10 arcsec resolution image with rms sensitivity of about $\sigma = 0.3$ mJy beam $^{-1}$ as input for TUNA. Fig. 8, left column, shows the comparison of the predicted mask with the 3σ contour of the LOFAR observation at a resolution of 80 arcsec with sensitivity $\sigma = 1$ mJy beam $^{-1}$. The confidence level for TUNA predictions is set to 0.9 to ensure reliability when the input resolution is lower than that used for training. Here, the input image has 10 arcsec resolution, while the network was trained on 6 arcsec resolution data.

Botteon et al. (2020), using LOFAR observations at 144 MHz, claimed the presence of a ~ 2 Mpc radio bridge connecting the pre-merging galaxy clusters A1758N and A1758S. Fig. 8, right column, shows the mask predicted by TUNA. The image at resolution ~ 6 arcsec with sensitivity $\sigma = 0.075$ mJy beam $^{-1}$ of the cluster PSZ2G107.10+65.32 from the LoTSS-DR2/PSZ2 data set has been used as input, whereas the 100 kpc tapered source-subtracted image at resolution ~ 30 arcsec with sensitivity $\sigma = 0.19$ mJy beam $^{-1}$ has been used as comparison.

TUNA shows the capability to effectively probe low surface brightness diffuse emission from radio bridges, underscoring its crucial role in detecting rare and elusive sources.

4.3 The megahaloes

Based on LOFAR observations, Cuciti et al. (2022) reported the existence of a new class of large diffuse radio sources that extend over 2–3 Mpc in four galaxy clusters, dubbed as radio megahaloes. The volume of the megahaloes is almost 30 times larger than that of radio haloes and their emissivity is about 20 times lower. We also used TUNA to detect these low surface brightness sources. The input data for Abell 665, Abell 697, and Abell 2218 have been cutout from the 20 arcsec resolution pointings provided by the LoTSS-DR2 survey (Shimwell et al. 2022b) (see Fig. 9). The ZwCl 0634.1+4750 cutout at 30 arcsec resolution has been obtained with the same data reduction and imaging procedures used for the survey (Cuciti et al. 2022).

For the megahaloes, TUNA trained at 20 arcsec was used. The results are presented in Fig. 10, where confidence contours are overlaid on the 120 arcsec resolution images. The network accurately captures both their size and morphology, confirming once more its ability to detect diffuse emissions at significantly lower resolutions than the input data.

5 CONCLUSIONS

Artificial intelligence offers novel solutions to achieve precise, automated, and rapid segmentation of astrophysical radio interferometric data. This paper demonstrates how vision transformers, integrated into the TransUNet architecture, enable new capabilities in radio astronomy image analysis, allowing for the detection of faint, diffuse radio sources that were inaccessible with state-of-the-art CNNs. We introduced TUNA, a deep learning model based on vision transformers, trained on cosmological simulations and applied without retraining to real observational data, thereby directly evaluating its

ability to generalize across domains. Our main achievements can be summarized as follows.

(i) TUNA detects extended radio sources at the sensitivity limits of the LOFAR HBA, without requiring manual source subtraction or low-resolution re-imaging. It can recover diffuse emission at scales equivalent to images reprocessed 4–6 times coarser than the original, while preserving morphological detail. Notable examples include the A399–A401 ridge and A1758 bridge, detected directly from high-resolution input, and all four known megahaloes, originally observed in low-resolution images produced via substantial UV tapering.

(ii) TUNA enables fully automated data processing, requiring no additional human supervision. The network can exploit HPC solutions for both training and inference, which significantly reduces the time to solution. For instance, The LoTSS-DR2 pointing P128+37, shown in Fig. 9, is processed in about 193 s at 6 arcsec (9528×9528 pixel) and in about 25 s at 20 arcsec (3176×3176 pixel), on Ampere100 NVIDIA GPUs of the Leonardo system at CINECA (see Section 3.3 for details). For comparison, the typical time needed to reduce one data set to the cluster field, subtract discrete sources and produce a low resolution image could require up to one day.

(iii) TUNA is highly effective in processing images with substantial noise and artefacts, recognizing only diffused sources. However, when lower resolution images are used as input, the risk of confusion noise increases and blended point sources may be interpreted by the network as diffused objects. The most effective strategy to mitigate this issue is to avoid reprocessing the data at coarser resolution and instead leverage the network’s ability to detect extended emission the native telescope resolution, thereby preserving source morphology and reducing misclassification.

(iv) The network succeeds at generalizing from simulated data to observational data. Moreover, it generalizes to diverse source types not present in the training set, including AGNs and their associated jets. This capacity opens up new opportunities for blind source detection and classification.

The methodology established with TUNA is undergoing further extension. The simulated training set is augmented by using AGN and point sources obtained from actual observations. This will provide a more accurate and thorough identification of radio galaxies and compact objects, distinguishing them from other extended sources. Future developments also include expanding its applicability across frequencies and instruments, and enhancing its output to estimate physical properties such as total flux densities alongside segmentation.

ACKNOWLEDGEMENTS

This paper is supported by the Fondazione ICSC, Spoke 3 Astrophysics, and Cosmos Observations. National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) Project ID CN.00000013” Italian Research Center for High-Performance Computing, Big Data and Quantum Computing’ funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di ‘campioni nazionali di R&S (M4C2-19)’ – Next Generation EU (NGEU), and it’s also supported by Programma Operativo Nazionale (PON), ‘Tematiche di Ricerca Green e dell’Innovazione’. We acknowledge the CINECA award under the IS CRA initiative, for the availability of high performance computing resources and support. The training and inference run this work is based on have been produced on the Leonardo Supercomputer at CINECA (Bologna, Italy) in the framework of the IS CRA

programme project IscrC-ViTBuild. We also acknowledge the usage of online storage tools kindly provided by the INAF Astronomical Archive (IA2) initiative (<http://www.ia2.inaf.it>). We thank Andrea Botteon for the set up of the LoTSS-DR2/PSZ2 archive and Chiara Stuardi, Bärbel Koribalski and Gabriele Giovannini for the insightful discussions.

DATA AVAILABILITY

Results from the TUNA processing of LoTSS-DR2/PSZ2 data are publicly available at:

(i) <https://owncloud.ia2.inaf.it/index.php/s/NtOMLBTXgJS2kZB>

The Sky and Clean images produced and adopted for the training and testing of the network are publicly available in FITS format at:

(i) <https://owncloud.ia2.inaf.it/index.php/s/IbFPICCCpUresrr>

The LOFAR radio images are publicly available at:

(i) <https://lofar-surveys.org/index.html>

REFERENCES

- Astropy Collaboration 2013, *A&A*, 558, A33
- Botteon A. et al., 2020, *MNRAS*, 499, L11
- Botteon A. et al., 2022, *A&A*, 660, A78
- Briggs D. S., 1995, American Astronomical Society Meeting Abstracts, 112.02
- Brown S. D., 2011, *JA&A*, 32, 577
- Brunetti G., 2011, *JA&A*, 32, 437
- Brunetti G., Jones T. W., 2014, *Int. J. Mod. Phys. D*, 23, 1430007
- Brunetti G., Vazza F., 2020, *Phys. Rev. Lett.*, 124, 051101
- Bryan G. L. et al., 2014, *ApJS*, 211, 19
- Cavanagh M. K., Bekki K., Groves B. A., 2021, *MNRAS*, 506, 659
- Chen J. et al., 2021, preprint ([arXiv:2102.04306](https://arxiv.org/abs/2102.04306))
- Cuciti V. et al., 2022, *Nature*, 609, 911
- DeepSeek-AI et al., 2024, preprint ([arXiv:2412.19437](https://arxiv.org/abs/2412.19437))
- Devlin J., Chang M.-W., Lee K., Toutanova K., 2019, in Burstein J., Doran C., Solorio T., eds, Proc. 2019 Conf. North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). Association for Computational Linguistics, Minneapolis, Minnesota, p. 4171, <https://aclanthology.org/19-1423/>
- Dosovitskiy A., 2020, preprint ([arXiv:2010.11929](https://arxiv.org/abs/2010.11929))
- Ferretti L., Giovannini G., Govoni F., Murgia M., 2012, *A&ARv*, 20, 1
- Gemini Team 2023, preprint ([arXiv:2312.11805](https://arxiv.org/abs/2312.11805))
- Gheller C., Vazza F., 2022, *MNRAS*, 509, 990
- Gheller C., Vazza F., Bonafede A., 2018, *MNRAS*, 480, 3749
- Govoni F. et al., 2019, *Science*, 364, 981
- He K., Zhang X., Ren S., Sun J., 2016, in Proc. IEEE Conf. Computer Vision and Pattern Recognition. IEEE, Las Vegas, p. 770
- Hodgson T., Vazza F., Johnston-Hollitt M., McKinley B., 2021, *Publ. Astron. Soc. Aust.*, 38, e047
- Hoefl M., Brügger M., 2007, *MNRAS*, 375, 77
- Hotan A. et al., 2021, *Publ. Astron. Soc. Aust.*, 38, e009
- Jonas J., Team M., 2016, MeerKAT Science: On the Pathway to the SKA. MeerKAT2016, Stellenbosch, p. 1
- Jones A. et al., 2023, *A&A*, 680, A31
- Kohavi R., 1995. Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI). Morgan Kaufmann Publishers Inc., San Francisco, p. 1137
- Lee W., Pillepich A., ZuHone J., Nelson D., Jee M. J., Nagai D., Finner K., 2024, *A&A*, 686, A55
- Nishiwaki K., Brunetti G., Vazza F., Gheller C., 2024, *ApJ*, 961, 15
- Offringa A. R., Smirnov O., 2017, *MNRAS*, 471, 301
- Offringa A. R., McKinley B., Hurley-Walker et al., 2014, *MNRAS*, 444, 606
- Pfrommer C., Springel V., Enßlin T. A., Jubelgas M., 2006, *MNRAS*, 367, 113
- Pignataro G. V. et al., 2024, *A&A*, 691, A99
- Radford A., Narasimhan K., Salimans T., Sutskever I., 2018, OpenAI Technical Report
- Radovic A. et al., 2018, *Nature*, 560, 41
- Rajpurohit K. et al., 2024, *ApJ*, 966, 38
- Rolnick D. et al., 2022, *ACM Comput. Surv. (CSUR)*, 55, 1
- Ronneberger O., Fischer P., Brox T., 2015, in Navab N., Hornegger J., Wells W. M., Frangi A., eds, Proc. 18th Int. Conf. Part 1: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015. Springer, Munich
- Ryu D., Kang H., Hallman E., Jones T. W., 2003, *ApJ*, 593, 599
- Ryu D., Kang H., Cho J., Das S., 2008, *Science*, 320, 909
- Sanvitale N., Gheller C., Bowman E., 2022, *Granular Matter*, 24, 65
- Shimwell T. et al., 2022a, *A&A*, 659, A1
- Shimwell T. W. et al., 2022b, *A&A*, 659, A1
- Stuardi C., Gheller C., Vazza F., Botteon A., 2024, *MNRAS*, 533, 3194
- Tingay S. J. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, e007
- Vamathevan J. et al., 2019, *Nature Rev. Drug Discov.*, 18, 463
- van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Van Weeren R., de Gasperin F., Akamatsu H., Brügger M., Ferretti L., Kang H., Stroe A., Zandanel F., 2019, *Space Sci. Rev.*, 215, 1
- Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A. N., Kaiser Ł., Polosukhin I., 2017, preprint ([arXiv:1706.03762](https://arxiv.org/abs/1706.03762)),
- Vazza F., Paoletti D., Banfi S., Finelli F., Gheller C., O’Sullivan S. P., Brügger M., 2021, *MNRAS*, 500, 5350
- Vazza F., Gheller C., Zanetti F., Tsizh M., Carretti E., Mtchedlidze S., Brügger M., 2025, *A&A*, 696, A58
- Wells D. C., Greisen E. W., Harten R. H., 1981, *A&AS*, 44, 363
- Wittor D., 2023, *Universe*, 9, 319

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: HYPERPARAMETERS OPTIMIZATION ANALYSIS

Table A1 presents the results of the evaluation of the impact of different hyperparameter settings on model performance. We varied tile size, batch size, learning rate, and Class 1 weight, and report the resulting Intersection over Union (IoU), precision, recall, and an average score combining the three metrics. The best 15 models are presented. The first line of the table shows our reference model. Changes in the hyperparameters result in small changes to the composite average score. Our tests indicate that a tile size of 512 and a learning rate between 0.001 and 0.005 yield optimal results. Learning rates smaller than 0.001 are absent from the initial 15 scores, due to lower performance. No clear trend was observed for batch size. Therefore, we adopted the value used in the original TransUnet implementation. The Class 1 weight w_1 is a multiplicative factor applied to the loss contribution of Class 1 pixels during training. It increases the contribution of underrepresented Class 1 pixels to the overall loss. In principle, this approach should improve model accuracy for the underrepresented class. In practice, however, it proved ineffective.

Table A1. Performance metrics (IoU, Precision, Recall, and composite Average Score) of the network across various hyperparameter configurations. Each row corresponds to a different combination of tile size, batch size, learning rate, and Class 1 weight w_1 . The average score is computed as the mean of IoU, precision, and recall to provide a single comparative value.

| Tile size | Batch size | Learning rate | w_1 | IoU | Precision | Recall | Avg. Score |
|-----------|------------|---------------|-------|------|-----------|--------|------------|
| 512 | 24 | 0.005 | 1 | 0.38 | 0.61 | 0.50 | 0.50 |
| 512 | 12 | 0.001 | 1 | 0.40 | 0.61 | 0.53 | 0.51 |
| 512 | 12 | 0.005 | 1 | 0.39 | 0.58 | 0.52 | 0.50 |
| 512 | 12 | 0.01 | 1 | 0.38 | 0.60 | 0.51 | 0.50 |
| 512 | 24 | 0.001 | 1 | 0.39 | 0.58 | 0.49 | 0.49 |
| 224 | 48 | 0.001 | 1 | 0.36 | 0.64 | 0.47 | 0.49 |
| 224 | 12 | 0.005 | 1 | 0.36 | 0.60 | 0.48 | 0.48 |
| 224 | 24 | 0.001 | 2 | 0.36 | 0.61 | 0.47 | 0.48 |
| 512 | 24 | 0.01 | 1 | 0.37 | 0.53 | 0.54 | 0.47 |
| 224 | 48 | 0.01 | 1 | 0.35 | 0.64 | 0.43 | 0.47 |
| 224 | 24 | 0.01 | 1 | 0.35 | 0.58 | 0.48 | 0.47 |
| 224 | 24 | 0.001 | 1 | 0.35 | 0.57 | 0.49 | 0.47 |
| 224 | 24 | 0.005 | 1 | 0.34 | 0.61 | 0.45 | 0.47 |
| 224 | 12 | 0.01 | 1 | 0.35 | 0.61 | 0.45 | 0.47 |
| 224 | 24 | 0.001 | 10 | 0.34 | 0.46 | 0.59 | 0.46 |

APPENDIX B: SUPPLEMENTARY TABLES

This appendix presents sample tables detailing the performance scores of TUNA (Table B1) and R-UNet (Table B2) as applied to the LoTSS-DR2/PSZ2 survey clusters at an average resolution

of 100 kpc, as well as the performance of TUNA on clusters at a resolution of 50 kpc (Table B3). The complete data tables, in tab-separated ASCII (TSV) format, are included as supplementary material.

Table B1. Performance scores of TUNA, considering the LoTSS-DR2/PSZ2 survey clusters ('Not Applicable' clusters excluded). The last column shows the average resolution of the 100 kpc source-subtracted tapered images used for evaluating the network's predictions. The average resolution is calculated as $(B_{\text{maj}} \times B_{\text{min}})^{1/2}$ where B_{maj} and B_{min} represent the major and minor axes of the beam, respectively. The full data table in ASCII format is available on-line as supplementary material.

| Cluster ID | Classification | Redshift | IoU | Precision | Recall | Accuracy | avg res 100 kpcSUB |
|-------------------|----------------|----------|-------|-----------|--------|----------|--------------------|
| PSZ2G023.17+86.71 | RH | 0.306 | 0.481 | 0.534 | 0.830 | 0.908 | 29 |
| PSZ2G031.93+78.71 | RH, U | 0.072 | 0.447 | 0.913 | 0.467 | 0.921 | 82 |
| PSZ2G040.58+77.12 | RH | 0.075 | 0.295 | 0.684 | 0.341 | 0.953 | 81 |
| PSZ2G045.13+67.78 | NDE | 0.219 | 0.168 | 0.450 | 0.211 | 0.963 | 45 |
| PSZ2G045.87+57.70 | RH | 0.611 | 0.404 | 0.430 | 0.871 | 0.921 | 20 |
| PSZ2G046.88+56.48 | RH | 0.115 | 0.182 | 0.856 | 0.188 | 0.818 | 56 |
| PSZ2G048.10+57.16 | RH, RR | 0.078 | 0.387 | 0.884 | 0.408 | 0.920 | 79 |
| PSZ2G048.75+53.18 | NDE | 0.098 | 0.102 | 0.188 | 0.183 | 0.970 | 64 |
| PSZ2G049.18+65.05 | NDE | 0.234 | 0.377 | 0.605 | 0.500 | 0.990 | 32 |
| PSZ2G049.32+44.37 | RH | 0.097 | 0.188 | 0.768 | 0.199 | 0.936 | 63 |

Table B2. Performance scores of R-UNet on the LoTSS-DR2/PSZ2 survey clusters ('Not Applicable' clusters excluded). The last column shows the average resolution of the 100 kpc source-subtracted tapered images used for evaluating the network's predictions. The average resolution is calculated as $(B_{\text{maj}} \times B_{\text{min}})^{1/2}$ where B_{maj} and B_{min} represent the major and minor axes of the beam, respectively. The full data table in ASCII format is available on-line as supplementary material.

| Cluster ID | Classification | Redshift | IoU | Precision | Recall | Accuracy | avg res 100 kpcSUB |
|-------------------|----------------|----------|-------|-----------|--------|----------|--------------------|
| PSZ2G023.17+86.71 | RH | 0.306 | 0.199 | 0.976 | 0.200 | 0.917 | 29 |
| PSZ2G031.93+78.71 | RH, U | 0.072 | 0.279 | 0.989 | 0.280 | 0.901 | 82 |
| PSZ2G040.58+77.12 | RH | 0.075 | 0.047 | 0.558 | 0.048 | 0.943 | 81 |
| PSZ2G045.13+67.78 | NDE | 0.219 | 0.000 | 0.000 | 0.000 | 0.965 | 45 |
| PSZ2G045.87+57.70 | RH | 0.611 | 0.275 | 0.627 | 0.328 | 0.947 | 20 |
| PSZ2G046.88+56.48 | RH | 0.115 | 0.051 | 0.890 | 0.051 | 0.794 | 56 |
| PSZ2G048.10+57.16 | RH, RR | 0.078 | 0.090 | 0.987 | 0.090 | 0.888 | 79 |
| PSZ2G048.75+53.18 | NDE | 0.098 | 0.000 | 0.000 | 0.000 | 0.977 | 64 |
| PSZ2G049.18+65.05 | NDE | 0.234 | 0.000 | 0.000 | 0.000 | 0.988 | 32 |
| PSZ2G049.32+44.37 | RH | 0.097 | 0.039 | 1.000 | 0.039 | 0.928 | 63 |

Table B3. Performance scores of TUNA considering the LoTSS-DR2/PSZ2 survey clusters ('Not Applicable' clusters excluded). The last column shows the average resolution of the 50 kpc source-subtracted tapered images used for evaluating the predictions of the network. The average resolution is calculated as $(B_{\text{maj}} \times B_{\text{min}})^{1/2}$ where B_{maj} and B_{min} represent the major and minor axes of the beam, respectively. The full data table in ASCII format is available on-line as supplementary material.

| Cluster ID | Classification | Redshift | IoU | Precision | Recall | Accuracy | avg res 50 kpcSUB |
|-------------------|----------------|----------|-------|-----------|--------|----------|-------------------|
| PSZ2G023.17+86.71 | RH | 0.306 | 0.409 | 0.426 | 0.912 | 0.901 | 21 |
| PSZ2G031.93+78.71 | RH, U | 0.072 | 0.572 | 0.775 | 0.686 | 0.959 | 43 |
| PSZ2G040.58+77.12 | RH | 0.075 | 0.331 | 0.395 | 0.674 | 0.977 | 41 |
| PSZ2G045.13+67.78 | NDE | 0.219 | 0.227 | 0.342 | 0.402 | 0.981 | 19 |
| PSZ2G045.87+57.70 | RH | 0.611 | 0.245 | 0.250 | 0.920 | 0.904 | 12 |
| PSZ2G046.88+56.48 | RH | 0.115 | 0.324 | 0.489 | 0.491 | 0.950 | 30 |
| PSZ2G048.10+57.16 | RH, RR | 0.078 | 0.501 | 0.676 | 0.660 | 0.962 | 40 |
| PSZ2G048.75+53.18 | NDE | 0.098 | 0.123 | 0.160 | 0.348 | 0.979 | 34 |
| PSZ2G049.18+65.05 | NDE | 0.234 | 0.307 | 0.533 | 0.420 | 0.988 | 22 |
| PSZ2G049.32+44.37 | RH | 0.097 | 0.256 | 0.419 | 0.398 | 0.977 | 33 |

This paper has been typeset from a \LaTeX file prepared by the author.