

2023-05

Morphological Classification of Radio Galaxies using Semi-Supervised Group Equivariant CNNs

Hossain, Mir Sazzat

Independent University, Bangladesh

<https://ar.iub.edu.bd/handle/11348/587>

Downloaded from IUB Academic Repository

Morphological Classification of Radio Galaxies using Semi-Supervised Group Equivariant CNNs

Mir Sazzat Hossain*, Sugandha Roy*, K. M. B. Asad*[†], Arshad Momen*[†], Amin Ahsan Ali*,
M Ashraful Amin* and A. K. M. Mahbubur Rahman*

*Center for Computational & Data Sciences, Independent University, Bangladesh

[†]Department of Physical Sciences, Independent University, Bangladesh

{mirsazzathossain, sugandha.roy08}@gmail.com, {kasad, arshad, aminali, aminmdashraful, akmmrahman}@iub.edu.bd

Abstract—Out of the estimated few trillion galaxies, only around a million have been detected through radio frequencies, and only a tiny fraction, approximately a thousand, have been manually classified. We have addressed this disparity between labelled and unlabelled images of radio galaxies by employing a semi-supervised learning approach to classify them into the known FRI and FRII types. A Group Equivariant Convolutional Neural Network was used as an encoder that preserves the equivariance for the Euclidean Group $E(2)$ to learn the representation of globally oriented feature maps through new Self-Supervised Learning (SSL) techniques SimCLR and BYOL. After representation learning, we trained a fully-connected classifier and fine-tuned the trained encoder with labelled data. We have found that this semi-supervised approach helps our method outperform a state-of-the-art method of classifying radio galaxies in many metrics. Our work reiterates the importance of semi-supervised learning in radio galaxy classification, where labelled data are scarce, but prospects are immense.

Index Terms—Radio Galaxy, Fanaroff-Riley, G-CNN, SimCLR, BYOL, Semi-supervised Learning

I. INTRODUCTION

Our universe is estimated to contain trillions of galaxies, but due to technological limitations, astronomers have only been able to observe a fraction of them. These observations are made at various frequencies such as optical (around 400–800 THz), infrared (around 1–30 THz) and radio (few MHz – few GHz). Galaxies appear to be different at different frequencies. Radio galaxies, which emit a significant amount of radiation at radio wavelengths, have been manually classified into various morphologies, including Fanaroff-Riley Type I and Type II [1], head-tailed [2], ringlike [3] and x-shaped [4].

The Fanaroff-Riley (FR) classification, based on the physical characteristics of a galaxy, is crucial for understanding the underlying science behind the formation and development of galaxies. Typical radio galaxies have a supermassive black hole in their core that actively accretes gas and stars along its equator and, as a result, ejects multiple jets through its poles. Consequently, radio images of a galaxy usually consist of a core and multiple lobes. FRI galaxies (Fig.1a) have the peak of their radio emission near the core, and the lobes have a darker edge than their core. On the other hand, FRII galaxies (Fig.1b) have peak emissions at the edge of the lobes far from the core.

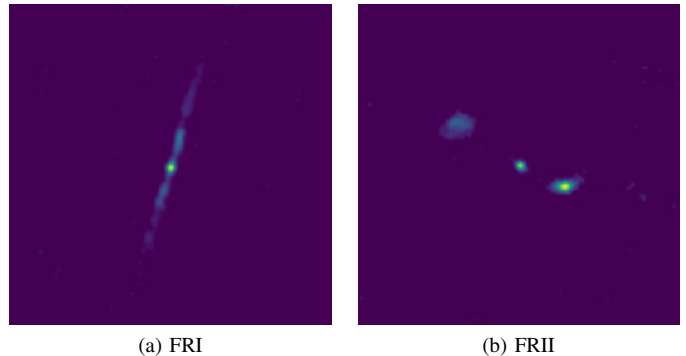


Fig. 1. Examples of Fanaroff-Riley galaxies. (a) FRI galaxy, characterized by peak radio emission near the core and darker edges of the lobes. (b) FRII galaxy, characterized by peak radio emission at the edge of the lobes far from the core.

Among almost a million detected radio galaxies, a mere fraction, approximately one thousand, have been manually classified as either FRI or FRII [5]. However, with the imminent operation of the Square Kilometre Array (SKA), the most significant scientific facility ever constructed, and its pathfinder telescopes, sensitivity and survey speed will increase significantly. This will enable the imaging of tens of thousands of radio sources in a matter of hours, a task that would have taken months to accomplish only a few years ago [6]. The sheer volume of data generated by the SKA requires using artificial intelligence to supplement human analysis. Hence, the radio astronomy and computing communities are actively engaged in the application of machine-learning techniques to classify radio galaxies.

The Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) survey, conducted using the Karl G. Jansky Very Large Array (VLA) radio telescope located in New Mexico, USA, provides some of the most detailed images of the radio sky at around 1.4 GHz. These images have been widely used in the classification of radio galaxies by researchers. Various techniques have been employed, such as using a Convolutional Neural Network (CNN) in the form of a slightly modified AlexNet as done by Aniyani and Thorat [7], or using a CNN with three convolution layers as demonstrated by Alhassan et al. [8]. Wu et al. [9] employed Faster Region-based Convo-

lutional Neural Networks (Faster R-CNN) to classify galaxies into six categories based on the number of components and peaks within a source. Tang et al. [10] adopted a 13-layer CNN and investigated the identification ability of classification networks on cross-survey data using Transfer Learning. Since deep networks provide better performance in classification, researchers are increasingly utilizing such networks to classify radio galaxies.

One of the main challenges in classifying radio galaxies using CNNs is the need for the network to be equivariant to different isometries such as translation, rotation, and mirror reflection. Since radio galaxies can have different orientations, CNN needs to be able to classify them regardless of their position or symmetry. While CNNs are typically translation-invariant, they must also be able to handle other types of isometries. Previous methods have attempted to address this issue by augmenting the data with rotated images. However, this approach has proven ineffective as the CNN may learn multiple copies of the same kernel differently. Scaife and Porter [11] addressed this issue by using Group Equivariant Convolutional Neural Networks (G-CNNs) [12], which can preserve equivariance for all isometries, allowing for more accurate classification of radio galaxies.

Another key challenge in classifying radio galaxies is utilising the vast amount of available unlabeled data. Self-supervised learning techniques can be used to extract useful information from these unlabeled data and apply it in a task-specific manner to improve the performance of our models. This approach effectively addresses the scarcity of labelled data in radio galaxy classification, as researchers have already begun to explore this method. For example, Ma et al. [13] developed a CNN-based autoencoder, MCRGNet, consisting of an encoder and a decoder block. First, they trained the MCRGNet autoencoder with unlabeled data and then fine-tuned the pre-trained encoder with labelled data to classify radio morphologies. This method performed significantly better than traditional supervised classifications [7], [8].

Ma et al. [14] used a VGG16 network as an autoencoder and achieved improved performance. However, more than a simple reconstruction error is needed to learn invariant representations under various noises and transformations. Slijepcevic et al. [15] used the semi-supervised learning technique FixMatch, which combines labelled and unlabeled data to improve model performance. This method uses a self-supervised approach to generate pseudo-labels for unlabeled data, which are then used to fine-tune the model with labelled data. This helps to extract useful information from large amounts of unlabeled data, resulting in improved model performance. However, the improvement is still relatively small compared to the baseline [11] with fewer labels [15] [16].

The advancement of recent self-supervised learning techniques, such as SimCLR (A Simple Framework for Contrastive Learning of Visual Representations) [17] and BYOL (Bootstrap Your Own Latent) [18] in representation learning and success of the “unsupervised pretrain, supervised fine-tune” paradigm compared to FixMatch for semi-supervised learning

[19] on the ImageNet dataset have motivated us to apply these techniques in the classification of radio galaxies. In our work, we have employed a semi-supervised approach using SimCLR and BYOL to utilize the underlying structure in the data for learning supervisory representations. These methods use losses that encourage the encoder to extract robust representations from large amounts of unlabeled data. To address the challenge of the need for the network to be equivariant to different isometries, such as translation, rotation, and mirror reflection, we have modified the encoders of the self-supervised models. Specifically, we have used a D16 (Dihedral group with 16 rotations) equivariant CNN proposed by Scaife and Porter [11] as a feature extractor in SimCLR and BYOL, which ensures that the extracted features are equivariant to the isometries. After representation learning, we fine-tuned the D16 equivalent CNN with labelled data and trained fully connected layers to classify radio galaxies. This approach effectively addresses the challenge of utilizing unlabeled data in radio galaxy classification. In summary, our paper presents the following key contributions:

- A semi-supervised approach for classifying radio galaxies utilizing self-supervised methods SimCLR, BYOL and G-CNN as encoder.
- Improved performance of our semi-supervised approach compared to state-of-the-art methods in FR classification.
- Evidence that representation learning from a large dataset of unlabeled data can be beneficial in identifying different types of galaxies, with potential application to tasks requiring less labelled data.

This paper is structured to provide a clear and organized analysis of our work on radio galaxy classification. In Section II, we discuss various models used in our research. Section III focuses on our proposed model and its unique training strategy. Section IV covers the datasets used in our study. Our experimental setup and results are presented in Sections V-A and V-B, respectively. Finally, in Section VI, we summarize our key findings and conclusions. We emphasize the superior performance of our radio galaxy classification model compared to other existing models.

II. BACKGROUND

A. Group Equivariant CNN (G-CNN)

Group Equivariant Convolutional Neural Network (G-CNN) [12] is a specialized neural network architecture that possesses the capability of being equivariant to symmetries present in the input data, such as translations, rotations, and reflections. This feature ensures that the network produces the same output when the input data is transformed in a particular way in accordance with the symmetries of a group.

G-CNN employs the theory of group representations to establish the convolution operation. This is achieved by substituting the traditional convolution operation with a group convolution operation, mathematically represented as

$$(f * \phi)(g) = \sum_{h \in G} f(h)\phi(h^{-1}g) \quad (1)$$

where f and ϕ are the input and kernel functions, respectively, G is the group of symmetries while g and h are its elements.

B. A Simple Framework for Contrastive Learning of Visual Representation (SimCLR)

SimCLR [17] is a self-supervised learning method that aims to learn representations of the input data by maximizing the agreement between two different views of the same data. The method utilizes the concept of contrastive learning, which is based on the idea of comparing the similarity between two different data samples.

SimCLR uses two neural networks, an encoder and a projection head. The encoder is trained to map the input data to a feature space, while the projection head is trained to map the features back to the input space. The two networks are trained together using a contrastive loss function, defined as the negative log-likelihood of the correct pair of views, mathematically

$$\mathcal{L}_{\text{SimCLR}} = -\log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{k=1}^{2n} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_k)/\tau)} \quad (2)$$

where z_i and z_j are the features of the two views of the same data sample, z_k are the features of a different sample; the dataset has n number of samples in total. $\text{sim}(\cdot)$ is a similarity function, such as cosine similarity, and τ is a temperature parameter. The objective of the loss function is to maximize the agreement between the two views of the same sample and minimize the agreement between different samples.

C. Bootstrap Your Own Latent (BYOL)

BYOL [18] is a self-supervised learning method that utilizes the idea of ‘‘global contrastive learning’’ to learn representations of input data. The method simultaneously trains two neural networks, an online network and a target network. The online network, parametrized by θ , is used to generate predictions of the input data. The target network, parametrized by ξ , is used to generate predictions of the predictions of the online network. The goal of BYOL is to learn similar representations between the two networks, which are, in turn, learned from the input data.

The training process of BYOL is guided by a contrastive loss function, defined as the mean squared error between the feature representations of the online network and the target network. Mathematically, this loss function is represented as

$$\mathcal{L}_{\theta, \xi} \triangleq \|\bar{q}_\theta(z_\theta) - \bar{z}'_\xi\|_2^2 = 2 - 2 \cdot \frac{\langle q_\theta(z_\theta), z'_\xi \rangle}{\|q_\theta(z_\theta)\|_2 \cdot \|z'_\xi\|_2} \quad (3)$$

where $q_\theta(z_\theta)$ is the feature representation of the input data generated by the online network, and z'_ξ is the feature representation of the input data generated by the target network. The goal of the loss function is to encourage the online network to learn representations similar to the representations of the target network.

The main advantage of BYOL is that it does not require explicit data augmentation or negative samples during training; it only requires the data to learn the representations.

III. PROPOSED METHOD

In this paper, we propose a semi-supervised approach for radio galaxy classification. Our method consists of two main steps: task-agnostic self-supervised learning and task-specific fine-tuning. In the first step, we use the state-of-the-art self-supervised representation learning techniques SimCLR and BYOL, to learn robust representations from a large amount of unlabeled data. In the second step, we fine-tune these representations using a small amount of labelled data for the specific task of radio galaxy classification. The use of a semi-supervised approach allows us to leverage a large amount of unlabeled data while still achieving high performance in task-specific classification. We have described these two steps in more detail and provided an in-depth analysis of our proposed method below.

A. Task-agnostic self-supervised learning

In this step, we employ the BYOL and SimCLR methods to learn representations of the input data using an unlabelled dataset. Fig. 2 illustrates the SimCLR and BYOL architectures adopted for self-supervised representation learning from the unlabeled dataset.

SimCLR (Fig. 2a) utilizes an encoder $f(\cdot)$ that extracts features from different augmented views x_i and x_j of the same image x . These augmented views generated from the same image are considered positive pairs of samples. Assuming a batch size of N , the other $2(N - 1)$ augmented examples are treated as negative pairs of samples. The representations, $h_i = f(x_i)$ and $h_j = f(x_j)$, are then mapped to the latent space using a projection head $g(\cdot)$, which is a Multi-Layer Perceptron (MLP) with one hidden layer. The contrastive loss is applied to the mapped representations $z_i = g(h_i)$ and $z_j = g(h_j)$.

BYOL (Fig. 2b) creates two augmented views, v and v' of an image. The feature extractor of the online network outputs a representation $y_\theta \triangleq f_\theta(v)$, and the projection head provides a projection $z_\theta \triangleq g_\theta(y_\theta)$ from the first augmented view v . From the second augmented view v' , the target network produces $y'_\xi \triangleq f'_\xi(v')$ and $z'_\xi \triangleq g'_\xi(y'_\xi)$. In the online network, an extra MLP is used to predict the output of the target encoder, i.e. $q_\theta(z_\theta)$ of z'_ξ . The loss function in BYOL is the mean squared error (MSE) between the l_2 -normalization of $q_\theta(z_\theta)$ and z'_ξ as shown in Eq. 3.

In the case of SimCLR and BYOL, the authors used a combination of cropping, colour distortion, and gaussian blur as data augmentation techniques. They found that this combination of augmentations improved the performance of these self-supervised models. However, the default augmentations used in SimCLR and BYOL were unsuitable for radio galaxy images because the galaxies occupy a very small portion of the images. Applying random cropping and resizing could result

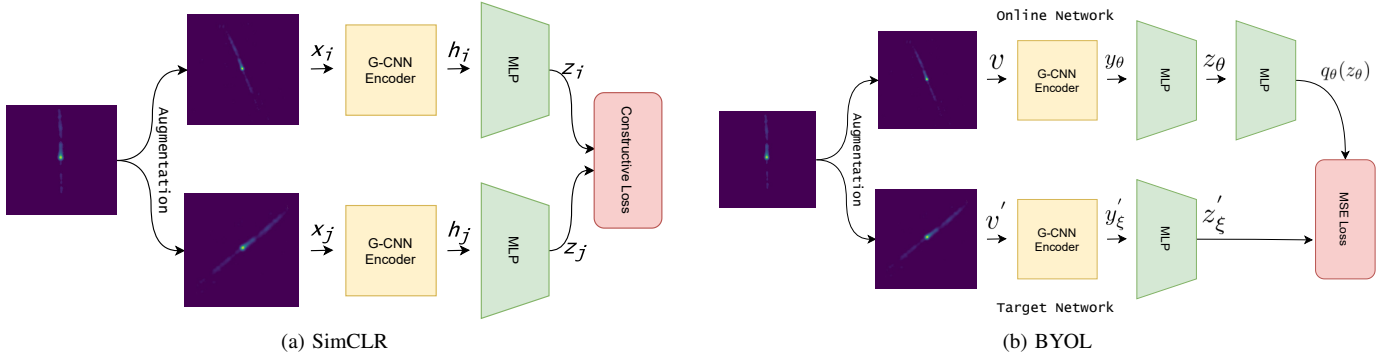


Fig. 2. Illustration of (a) SimCLR and (b) BYOL architectures used for self-supervised representation learning from an unlabeled dataset. The data augmentation strategy for radio galaxy images has been modified by setting bounds for aspect ratio, adding an extra random rotation, and replacing the ResNet-50 feature extractor with an E(2)-Equivariant Steerable G-CNN for improved performance on the downstream radio galaxy classification task.

in many empty patches which would negatively impact the model’s training.

Our contribution to this work is redesigning the data augmentation policy for radio galaxy images. We retained the colour distortion and Gaussian blur augmentations, set lower and upper bounds for the random aspect ratio of the crop to 0.8 and 1, respectively, and added an extra 360° random rotation to the augmentations. This helped to ensure that most of the cropped patches contained the galaxy while increasing the diversity of the training data. By redesigning the data augmentation strategy, we improved the quality of the learned representations and achieved better results.

In addition to redesigning the data augmentation strategy, we have also made another significant contribution to this work by replacing the ResNet-50 feature extractor of both BYOL and SimCLR with the E(2)-Equivariant Steerable G-CNN, specifically the D16 Steerable G-CNN. This network contains 2 group convolution layers with kernel size 5×5 , each followed by a relu activation and max pooling. The output is then passed through a linear layer that maps the 20736-dimensional feature maps to a feature map of dimension 2048. This feature map is then used as input for the MLP projection head. This change allows the network to be equivariant to different isometries, enabling it to classify radio galaxies with different orientations, regardless of their position or symmetry.

B. Task-specific fine-tuning

In order to optimize the model for the task of FR classification, we perform task-specific finetuning on our labelled dataset next. This step involves modifying the pre-trained encoder obtained through task-agnostic self-supervised learning to suit better our labelled data’s characteristics.

As depicted in Fig. 3, the final fully connected layer of the encoder is replaced with a new architecture that consists of three fully connected linear layers. The first linear layer transforms the 20736-dimensional feature maps into a 120-dimensional feature map, the second linear layer maps the 120-dimensional feature maps to an 84-dimensional feature map, and finally, the third linear layer reduces it to a 2-dimensional feature map. We applied ReLU activation functions after the

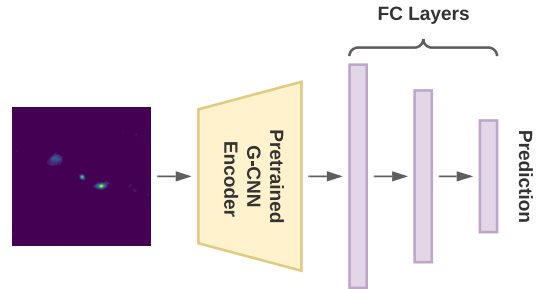


Fig. 3. Illustration of the finetuning step for FR classification. The pre-trained encoder from task-agnostic self-supervised learning is modified by replacing the last fully connected layer with three new fully connected layers, allowing the model to adapt to the specific task of FR classification.

first and second linear layers and a dropout layer after the second linear layer.

This modification allows the network to adapt to the specific task of FR classification and enhances its ability to recognize the distinctive characteristics of the labelled data. The finetuning is done using a supervised learning approach and cross-entropy loss function on the labelled data. This step significantly improves the model’s performance on the task of FR classification by adapting to the unique characteristics of our labelled data.

The combination of our task-agnostic self-supervised learning approach and task-specific fine-tuning provides a powerful solution for radio galaxy classification. Our proposed method leverages vast amounts of high-quality images. It fine-tunes the model to the unique characteristics of our labelled data, resulting in improved accuracy and faster convergence compared to traditional supervised learning.

IV. DATASETS

Throughout this work, we utilized radio galaxy images obtained from the Karl G. Jansky Very Large Array (VLA) as part of the Faint Images of the Radio Sky at Twenty-Centimeters (FIRST) survey. The survey covered a vast area of 10,000 square degrees of the sky, producing images with a

TABLE I
NUMERICAL IDENTIFIERS OF MIRABEST BATCHED DATASET

Digit 1	Digit 2	Digit 3
1 - FRI	0 - Confident	0 - Standard
2 - FRII	1 - Uncertain	1 - Double-double
3 - Hybrid		2 - Wide-angle Tail
		3 - Diffuse
		4 - Head-tail

resolution of 5 arcsec and a sensitivity of 0.15 mJy (milli-jansky; $1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$). Each pixel in the final images corresponds to an angular size of 1.8 arcsec. We utilized pre-processed images from the FIRST survey to perform representation learning and radio galaxy classification.

For representation learning, we used the Radio Galaxy Zoo (RGZ) dataset [20], pre-processed by Wu et al. [9]¹. The original dataset consisted of 11,678 images², of which we selected approximately 9,700 images where only one galaxy was visible. Each image originally had dimensions of 132×132 pixels, subsequently padded with zeros to 150×150 pixels. This unlabelled dataset will be referred to as Dataset-U in the following sections.

For the fine-tuning step with labelled dataset, we have used FIRST images labelled by Miraghaei and Best [5] and pre-processed for FR classification by Porter [21]. The dataset, referred to as the MiraBest Batched Dataset, contains 1256 images each of which is assigned a three-digit numerical identifier. The first digit of this identifier denotes the source class: 1 for FRI, 2 for FRII, and 3 for the hybrid sources. The second digit indicates the level of confidence in the classification: 0 for confidently classified and 1 for uncertainly classified. The third digit denotes the morphology of the source: 0 for standard morphology, 1 for double-double, 2 for wide-angled tail, 3 for diffuse source, and 4 for head-tail. The meaning of each digit is summarized in Table I.

In our work, we have only utilized confident sources from the MiraBest Batched Dataset, specifically those with numerical identifications 100, 102, 104, 200, or 201. We named this labelled data Dataset-F which was utilized for training and testing purposes. Like Dataset-U, each image in Dataset-F has a resolution of 150×150 pixels. The breakdown of source counts for FRI and FRII sources in the training and testing sets of Dataset-F can be found in Table II.

V. EXPERIMENTAL ANALYSIS

This section presents our experimental setup and provides a detailed analysis of the results obtained through our semi-supervised approach for radio galaxy classification.

A. Experimental setup

In order to evaluate the effectiveness of our proposed approach, we conducted experiments using two self-supervised

TABLE II
SOURCE COUNTS FOR FRI AND FRII SOURCES IN THE LABELLED DATASET (DATASET-F)

Morphology	MiraBest		Total
	Train	Test	
FRI	348	49	397
FRII	381	55	436
Total	729	104	833

methods, SimCLR and BYOL, separately on the unlabeled dataset, Dataset-U, to learn the underlying representations of the radio galaxy images.

For the SimCLR model, we fixed the following hyperparameters: (i) learning rate of 0.02, (ii) weight decay of 10^{-6} , (iii) batch size of 16, (iv) MLP hidden layer dimension of 2048, (v) MLP output layer dimension of 128, and (vi) number of epochs of 500. We used the Layer-wise Adaptive Rate Scaling (LARS) optimizer and Cosine Annealing Warm Restarts scheduler with a minimum learning rate of 0.05.

For the BYOL model, we tuned the following hyperparameters: (i) learning rate of 3×10^{-4} , (ii) moving average decay of 0.99, (iii) batch size of 16, (iv) MLP hidden layer dimension of 4092, (v) MLP output layer dimension of 256, and (vi) number of epochs of 500. We used the Adam optimizer to train this model.

Once the representations were learned from the unlabeled dataset, we used the learned representations to fine-tune the G-CNN model using a supervised approach on the labelled dataset, Dataset-F. We tuned the following hyperparameters for this step: (i) learning rate of 0.0001, (ii) weight decay of 10^{-6} , (iii) batch size of 50, and (iv) number of epochs of 600. We used Adam optimizer and a Reduce Learning Rate (LR) on Plateau strategy with patients of 2 and a factor of 0.9.

The experiments in this paper were carried out on a high-performance computing environment equipped with a powerful Intel(R) Core(TM) i9-9900K CPU @ 3.60GHz and 64 Gigabytes of RAM, paired with a cutting-edge GeForce RTX 2080 Ti GPU. The training process for each self-supervised model took approximately 17 hours to complete on this setup, while the downstream training required roughly 40 minutes to run.

B. Result Analysis

In this section, we present a comprehensive analysis of our proposed method for classifying radio galaxies. We have conducted several experiments to compare our approach with a traditional supervised approach and demonstrate our model's effectiveness using different qualitative and quantitative metrics.

At first, we evaluated the quality of the clusters formed by learned representations with our semi-supervised approaches (BYOL and SimCLR) and compared the results with the supervised model. To do so, we passed the labeled dataset (Dataset-F) through the finetuned encoder for semi-supervised approaches (BYOL, SimCLR). We extracted the representation from the first fully connected (FC) layer. We have performed

¹https://github.com/chenwuperth/rgz_rcnn

²<https://cloudstor.aarnet.edu.au/plus/s/agKNekOJK87hOh0>

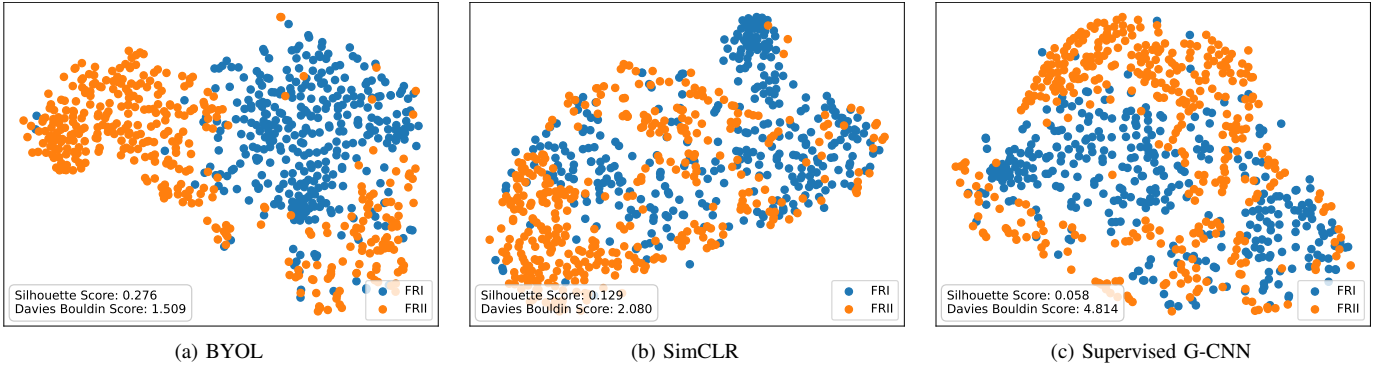


Fig. 4. Visualization of the representations obtained from the first fully connected layer of the fine-tuned encoders of (a) BYOL, (b) SimCLR, and (c) the Supervised G-CNN Model on our labelled dataset using t-SNE. The blue points represent the FRI class, while the orange points represent the FR II class. Our models, as evidenced by the higher Silhouette scores and the lower Davies Bouldin scores, show improved clustering performance, demonstrating that the models learned more effective data representations than the supervised model.

the same process for the supervised model. The output of the first FC layer is a 120-dimensional vector, which is impossible to visualize. We employed t-Distributed Stochastic Neighbor Embedding (t-SNE) to visualize the feature maps in a lower-dimensional space. The results of the t-SNE visualization, plotted in Fig. 4, indicate that the cluster quality formed by the representations learned from the semi-supervised models is significantly better than those from the supervised model. It means the data points are more clearly clustered.

In order to further validate our observations, we used two widely accepted metrics to quantify the cluster quality: Silhouette Score and Davies Bouldin Score. The Silhouette Score measures the similarity of a sample to its own cluster compared to other clusters, with higher scores indicating better cluster formation. On the other hand, the Davies Bouldin Score measures the average similarity between each cluster and its closest cluster, with lower scores representing better cluster quality. The results of our evaluation showed that the fine-tuned encoders of BYOL and SimCLR had Silhouette Scores of 0.276 and 0.129, respectively, while the supervised model had a score of 0.058. The BYOL and SimCLR models also received Davies Bouldin Scores of 1.509 and 2.080, respectively, compared to 4.814 for the supervised model. These results demonstrate that our semi-supervised models have learned more effective data representations, resulting in significantly improved clustering performance when compared to the supervised model.

In order to evaluate the convergence of our models during the fine-tuning stage, we performed 5-fold cross-validation on the labelled dataset (Dataset-F) using the representations learned by the encoders of SimCLR and BYOL. We also compared the convergence of the loss function in the fine-tuning step with that of the supervised G-CNN approach. For each repetition, we used a different fold for validation and the remaining four folds for training. Fig. 5a and 5b show the convergence plots for the training and validation loss, respectively. The plots were generated by taking the mean and standard deviation of the loss across the five repeats. As seen from

the plots, the training and validation losses for our fine-tuning approach converge very quickly compared to the supervised G-CNN approach, indicating that the representations learned in the first stage are adequate for the downstream task.

In order to further evaluate the performance of our semi-supervised models, we fine-tuned the encoders of SimCLR and BYOL using a labelled dataset (Dataset-F) and compared the results to a state-of-the-art supervised model, D16-Steerable G-CNN. We trained and tested all models using the same setup and repeated the process 15 times for each model. As shown in Table III, the results include the average precision, recall, and F1 scores for both FRI and FR II classes, along with the standard deviation. The best metrics are highlighted in bold, and the second-best is underlined.

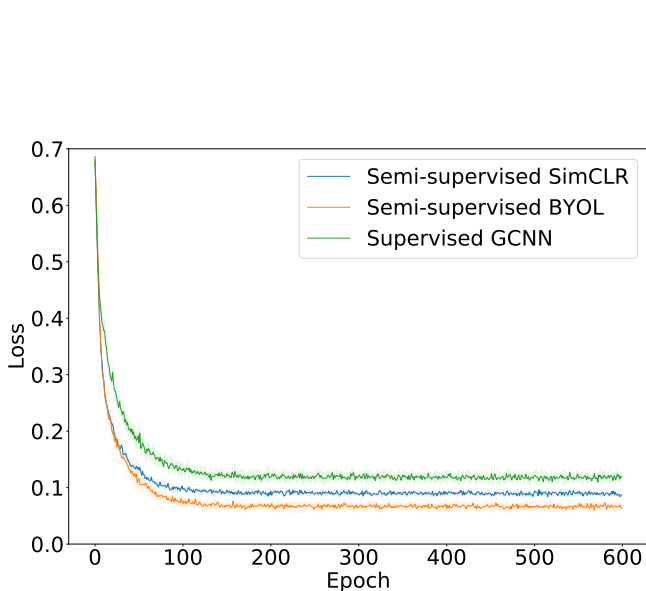
The results demonstrate that our semi-supervised models outperform the supervised model. Both BYOL and SimCLR achieved higher accuracy than D16-Steerable G-CNN. BYOL achieved the highest accuracy of 97.12%, followed by SimCLR at 95.77%. In addition, both models exhibited superior performance in other classification metrics. For example, SimCLR achieved precision and recall of 98% for the FRI and FR II radio galaxies, respectively, and BYOL achieved an excellent f-score of 97% for both classes. These results demonstrate the effectiveness of our proposed semi-supervised models in accurately classifying radio galaxies.

In addition to the results obtained from the classification metrics, we further evaluated the performance of our models using Receiver Operating Characteristic (ROC) curves and the corresponding Area Under the Curve (AUC) scores. The ROC curve is a graphical representation of the performance of a binary classifier system, while the AUC score provides a single numeric value to represent the classifier’s performance. It ranges from 0 to 1, with a score of 1 indicating a perfect classifier and a score of 0.5 indicating a classifier that performs no better than random.

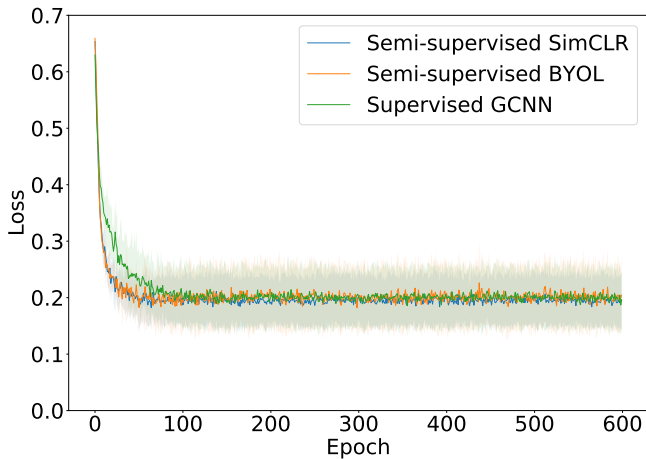
The evaluation results are depicted in Fig. 6. The AUC scores for BYOL and SimCLR’s fine-tuned encoders were found to be 0.99 and 0.98, respectively, while the score for the

TABLE III
PERFORMANCE COMPARISON BETWEEN SEMI-SUPERVISED AND SUPERVISED METHODS; THE TABLE ILLUSTRATES THE SUPERIORITY OF OUR SEMI-SUPERVISED MODELS OVER THE STATE-OF-THE-ART SUPERVISED METHOD ACROSS VARIOUS CLASSIFICATION METRICS

	Accuracy[%]	FRI			FRII		
		Precision	Recall	f1-score	Precision	Recall	f1-score
Semi-supervised SimCLR	95.77 ± 0.90	0.98 ± 0.061	0.93 ± 0.018	0.95 ± 0.011	0.94 ± 0.013	0.98 ± 0.014	0.96 ± 0.009
Semi-supervised BYOL	97.12 ± 0.40	0.97 ± 0.008	<u>0.96 ± 0.009</u>	0.97 ± 0.005	<u>0.96 ± 0.007</u>	<u>0.98 ± 0.008</u>	0.97 ± 0.004
Supervised G-CNN	94.80 ± 0.90	0.93 ± 0.012	0.96 ± 0.010	0.94 ± 0.009	0.96 ± 0.009	0.94 ± 0.012	0.95 ± 0.009



(a) Training Loss



(b) Validation Loss

Fig. 5. Convergence plots for fine-tuning the encoders of SimCLR and BYOL, compared to training a supervised G-CNN on Dataset-F. The plots show the mean and standard deviation of (a) training loss and (b) validation loss across 5-fold cross-validation. The fast convergence of the fine-tuned encoders compared to the supervised G-CNN suggests the effectiveness of the learned representations.

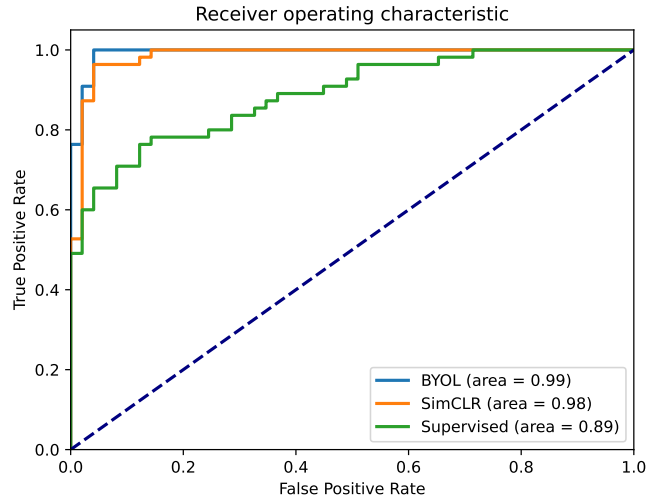


Fig. 6. Receiver Operating Characteristic (ROC) curves for the fine-tuned encoders of BYOL, SimCLR and the Supervised model, displaying the performance of the models in classifying the two classes of radio galaxies. The corresponding Area Under the Curve (AUC) scores are also shown, with a score of 1 indicating a perfect classifier and a score of 0.5 indicating a classifier that performs no better than random.

Supervised model was 0.89. These high AUC scores for BYOL and SimCLR indicate their exceptional ability to distinguish between the two classes of radio galaxies accurately. The models can effectively differentiate between radio galaxies with a minimal rate of false positives, which implies that they can correctly identify most radio galaxies without misclassifying them.

Overall, these results support the effectiveness of our proposed semi-supervised models in classifying radio galaxies. The high AUC scores and the shape of the ROC curves indicate that our models can effectively differentiate between the two classes of radio galaxies. This makes them valuable for further analysis and studies in the field.

To further demonstrate the superiority of our proposed models over the supervised approach, we conducted a statistical t-test on the accuracy scores of 15 runs of each model. The paired t-test between the supervised G-CNN and SimCLR models yielded a t-value of approximately -1.89 and a p-value of approximately 0.08. Similarly, the test between the supervised G-CNN and BYOL models yielded a t-value of approximately -3.47 and a p-value of approximately 0.0038. These results indicate that our proposed models perform

significantly better than the supervised G-CNN, with a high level of statistical significance.

VI. CONCLUSION

Our proposed semi-supervised group-equivariant CNN approach has been proven to be an effective solution for classifying radio galaxies. The use of a vast amount of unlabeled images during the pre-training stage allowed our models to learn robust and invariant features that can withstand noise and transformations present in radio galaxy images. The equivariant group architecture further enhanced the model's ability to extract structural and symmetry information, which is crucial for accurate classification.

Fine-tuning the pre-trained encoder with a limited amount of labelled data significantly improved its performance compared to traditional supervised methods. Our quantitative analysis, using Silhouette Score and Davies Bouldin Score, showed that the fine-tuned encoders produced higher Silhouette scores and lower Davies Bouldin scores, indicating better cluster formation compared to the supervised model. The visualization of the representations, obtained from the first fully connected layer of the fine-tuned encoders, demonstrated a clear separation between FRI and FRII galaxies into different clusters (Fig. 4). This highlights the effectiveness of our task-agnostic self-supervised learning approach in utilizing limited labelled data to learn effective data representations.

Our semi-supervised method significantly improved over traditional supervised methods, achieving a high accuracy of 97.12% using the fine-tuned BYOL encoder compared to only 94.80% for the supervised approach (Table III). Furthermore, our training and validation losses converged much faster than in the supervised approach (Fig. 5), and the ROC curves and AUC scores (Fig. 6) confirmed the ability of our models to distinguish between the two classes of radio galaxies effectively.

To further solidify our findings, a paired t-test on the accuracy score of 15 runs showed a statistically significant improvement in our proposed models over the supervised model. These results highlight the immense potential of semi-supervised learning in classifying radio galaxies and demonstrate the importance of utilizing vast amounts of high-quality images.

Our work is a stepping stone towards utilizing this powerful technique for classifying the millions of radio galaxies yet to be discovered and can lead to further research aimed at improving classification accuracy. Not only for radio galaxies but our proposed semi-supervised group-equivariant CNN approach also offers a promising solution for other classification tasks that have limited labelled data. With upcoming radio telescopes such as SKA, the amount of high-quality images available for astronomical object classification is set to increase. Our approach provides a scalable and efficient solution to meet this challenge.

VII. ACKNOWLEDGMENT

This project has been jointly sponsored by Independent University, Bangladesh and the ICT Division of the Bangladesh Government.

REFERENCES

- [1] B. L. Fanaroff and J. M. Riley, "The morphology of extragalactic radio sources of high and low luminosity," *Monthly Notices of the Royal Astronomical Society*, vol. 167, pp. 31P–36P, May 1974.
- [2] T. K. Sasmal, S. Bera, S. Pal, and S. Mondal, "A New Catalog of Head-Tail Radio Galaxies from the VLA FIRST Survey," *Astrophysical Journal, Supplement*, vol. 259, no. 2, p. 31, Apr. 2022.
- [3] D. D. Proctor, "Morphological Annotations for Groups in the First Database," *Astrophysical Journal, Supplement*, vol. 194, no. 2, p. 31, Jun. 2011.
- [4] J. P. Leahy and P. Parma, "Multiple outbursts in radio galaxies," in *Extragalactic Radio Sources. From Beams to Jets*, J. Roland, H. Sol, and G. Pelletier, Eds., Jan. 1992, pp. 307–308.
- [5] H. Miraghaei and P. N. Best, "The nuclear properties and extended morphologies of powerful radio galaxies: the roles of host galaxy and environment," *Monthly Notices of the Royal Astronomical Society*, vol. 466, no. 4, pp. 4346–4363, Apr. 2017.
- [6] T. J. Galvin, M. T. Huynh, R. P. Norris, X. R. Wang, E. Hopkins, K. Polsterer, N. O. Ralph, A. N. O'Brien, and G. H. Heald, "Cataloguing the radio-sky with unsupervised machine learning: a new approach for the SKA era," *Monthly Notices of the Royal Astronomical Society*, vol. 497, no. 3, pp. 2730–2758, Sep. 2020.
- [7] A. K. Aniyani and K. Thorat, "Classifying Radio Galaxies with the Convolutional Neural Network," *Astrophysical Journal, Supplement*, vol. 230, no. 2, p. 20, Jun. 2017.
- [8] W. Alhassan, A. R. Taylor, and M. Vaccari, "The FIRST Classifier: compact and extended radio galaxy classification using deep Convolutional Neural Networks," *Monthly Notices of the Royal Astronomical Society*, vol. 480, no. 2, pp. 2085–2093, Oct. 2018.
- [9] C. Wu, O. I. Wong, L. Rudnick, S. S. Shabala, M. J. Alger, J. K. Banfield, C. S. Ong, S. V. White, A. F. Garon, R. P. Norris, H. Andernach, J. Tate, V. Lukic, H. Tang, K. Schawinski, and F. I. Diakogiannis, "Radio Galaxy Zoo: CLARAN - a deep learning classifier for radio morphologies," *Monthly Notices of the Royal Astronomical Society*, vol. 482, no. 1, pp. 1211–1230, Jan. 2019.
- [10] H. Tang, A. M. M. Scaife, and J. P. Leahy, "Transfer learning for radio galaxy classification," *Monthly Notices of the Royal Astronomical Society*, vol. 488, no. 3, pp. 3358–3375, Sep. 2019.
- [11] A. M. M. Scaife and F. Porter, "Fanaroff-Riley classification of radio galaxies using group-equivariant convolutional neural networks," *Monthly Notices of the Royal Astronomical Society*, vol. 503, no. 2, pp. 2369–2379, May 2021.
- [12] T. Cohen and M. Welling, "Group equivariant convolutional networks," in *International conference on machine learning*. PMLR, 2016, pp. 2990–2999.
- [13] Z. Ma, H. Xu, J. Zhu, D. Hu, W. Li, C. Shan, Z. Zhu, L. Gu, J. Li, C. Liu, and X. Wu, "A Machine Learning Based Morphological Classification of 14,245 Radio AGNs Selected from the Best-Heckman Sample," *Astrophysical Journal, Supplement*, vol. 240, no. 2, p. 34, Feb. 2019.
- [14] Z. Ma, J. Zhu, and H. Zhu, Yongkai Xu, "Classification of Radio Galaxy Images with Semi-supervised Learning," in *Data Mining and Big Data*, ser. Communications in Computer and Information Science, Tan, Ying and Shi, Yuhui, Ed., vol. 1071. 4th International Conference, DMBD 2019, Chiang Mai, Thailand: Springer, July 2019, pp. 191–200.
- [15] I. V. Slijepcevic, A. M. M. Scaife, M. Walmsley, M. Bowles, O. I. Wong, S. S. Shabala, and H. Tang, "Radio Galaxy Zoo: using semi-supervised learning to leverage large unlabelled data sets for radio galaxy classification under data set shift," *Monthly Notices of the Royal Astronomical Society*, vol. 514, no. 2, pp. 2599–2613, Aug. 2022.
- [16] I. V. Slijepcevic and A. M. M. Scaife, "Can semi-supervised learning reduce the amount of manual labelling required for effective radio galaxy morphology classification?" *arXiv e-prints*, p. arXiv:2111.04357, Nov. 2021.
- [17] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.

- [18] J.-B. Grill, F. Strub, F. Alché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. D. Guo, M. Gheshlaghi Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko, “Bootstrap your own latent: A new approach to self-supervised Learning,” *arXiv e-prints*, p. arXiv:2006.07733, Jun. 2020.
- [19] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. Hinton, “Big self-supervised models are strong semi-supervised learners,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS’20. Red Hook, NY, USA: Curran Associates Inc., 2020.
- [20] J. K. Banfield, O. I. Wong, K. W. Willett, R. P. Norris, L. Rudnick, S. S. Shabala, B. D. Simmons, C. Snyder, A. Garon, N. Seymour, E. Middelberg, H. Andernach, C. J. Lintott, K. Jacob, A. D. Kapińska, M. Y. Mao, K. L. Masters, M. J. Jarvis, K. Schawinski, E. Paget, R. Simpson, H. R. Klöckner, S. Bamford, T. Burchell, K. E. Chow, G. Cotter, L. Fortson, I. Heywood, T. W. Jones, S. Kaviraj, Á. R. López-Sánchez, W. P. Maksym, K. Polsterer, K. Borden, R. P. Hollow, and L. Whyte, “Radio Galaxy Zoo: host galaxies and radio morphologies derived from visual inspection,” *Monthly Notices of the Royal Astronomical Society*, vol. 453, no. 3, pp. 2326–2340, Nov. 2015.
- [21] F. A. M. Porter, “Mirabest batched dataset,” Nov. 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.4288837>