

# Radio Galaxy Classification with wGAN-Supported Augmentation

Janis Kummer<sup>1,2</sup>, Lennart Rustige<sup>1,3</sup>, Florian Grieser<sup>1,4</sup>, Kerstin Borrás<sup>3,5</sup>,  
Marcus Brüggemann<sup>2</sup>, Patrick L. S. Connor<sup>1,6</sup>, Frank Gaede<sup>3</sup>, Gregor Kasieczka<sup>6</sup>,  
Peter Schleper<sup>6</sup>

**Abstract:** Novel techniques are indispensable to process the flood of data from the new generation of radio telescopes. In particular, the classification of astronomical sources in images is challenging. Morphological classification of radio galaxies could be automated with deep learning models that require large sets of labelled training data. Here, we demonstrate the use of generative models, specifically Wasserstein GANs (wGAN), to generate artificial data for different classes of radio galaxies. Subsequently, we augment the training data with images from our wGAN. We find that a simple fully-connected neural network for classification can be improved significantly by including generated images into the training set.

**Keywords:** Radio galaxy classification, Generative models, GANpification

## 1 Introduction

The new generation of radio telescopes (e.g. LOFAR, MeerKAT and in the future the SKA [Ca04; JM16; va13]) will produce enormous amounts of data and the improved sensitivity of the instruments leads to a much higher source density within these data sets. As a result, a new level of automation for processing the data and for classifying sources is needed. Morphological classification of radio sources is crucial for answering a range of fundamental astrophysical questions, such as the origin of cosmic magnetism. One promising approach for morphological classification relies on the use of deep classifiers trained on well-understood data sets. However, the existing amount of data with morphological labels is limited, as they are typically extracted from catalogues created and curated manually by experts. Small data

---

<sup>1</sup> Center for Data and Computing in Natural Sciences (CDCS), Notkestrasse 9, D-22607 Hamburg, Germany

<sup>2</sup> Universität Hamburg, Hamburger Sternwarte, Gojenbergsweg 112, D-21029 Hamburg, Germany  
janis.kummer@uni-hamburg.de

<sup>3</sup> Deutsches Elektronen-Synchrotron DESY, Notkestrasse 85, D-22607 Hamburg, Germany  
lennart.rustige@desy.de

<sup>4</sup> Section for Biomedical Imaging, University Medical Center Hamburg-Eppendorf, D-20246 Hamburg, Germany  
Institute for Biomedical Imaging, Hamburg University of Technology, D-21073 Hamburg, Germany  
florian.grieser@tuhh.de

<sup>5</sup> RWTH Aachen University, Templergraben 55, D-52062 Aachen, Germany

<sup>6</sup> Universität Hamburg, Institut für Experimentalphysik, Luruper Chaussee 149, D-22761 Hamburg, Germany

sets used in the training of deep learning models for galaxy classification can be enlarged by data augmentation, e.g. by applying random rotations and reflections to the images (classical augmentation). In this work, we investigate a novel application of generative models to enhance the available training sets. For this augmentation technique, multiple neural networks are combined to learn the underlying distribution of a data set. This allows the creation of new data by sampling from the learnt distribution. In this study we investigate whether radio galaxy classifiers can be improved when trained on such enhanced data sets containing generated data. For similar approaches from different fields see for example [Fr18; Zh17].

As radio sources exhibit a large variety of structures, we consider a four-class classification problem (as in [ATV18; Sa21]), including bent-tail and compact sources in addition to the classes FRI and FRII of Fanaroff; Riley [FR74]. For the class FRI, the maximum of the radio emission is situated close to the centre of the source. The maxima of the radio emission are located at the edges of the jets for FRII sources. Unresolved and point sources are contained in the Compact class. The Bent class consists of sources for which the angle between the jets differs significantly from 180 degrees. The two subtypes narrow-angle tail (NAT) and wide-angle tail (WAT) are further discriminated by the angle. For an illustration of the considered classes (FRI, FRII, Compact and Bent) see Fig. 1.




Fig. 1: Examples of the classes FRI, FRII, Compact and Bent.

Aniyan; Thorat [AT17], Alhassan et al. [ATV18], Samudre et al. [Sa21], and Tang et al. [TSL19] use convolutional neural networks (CNNs) trained on data from the FIRST survey [BWH95] for the classification of radio galaxies.

## 2 Data

We combine different catalogues that characterise radio sources from the FIRST survey to create a data set of radio galaxy images with morphological labels [BCM17; CMB17a; CMB17b; GBW10; GW08; MB17; Pr11]. By explicitly comparing the coordinates of the sources, 300 duplicates were found and removed. Moreover, we exclude 147 sources from the data set that appear in different catalogues with different labels. The resulting data sets are presented in Table 1. We adopt the preprocessing from [AT17] after cropping the images to the input size of our generative network (128 x 128 pixels). In particular, we set all pixel values below three times the local RMS noise to the value of this threshold. Subsequently, the pixel values are rescaled to the range between -1 and 1 to represent floating point greyscale images. We apply classical augmentation to each image during training.

	FRI	FRII	Compact	Bent	Total
train	395	824	291	248	1758
validation	50	50	50	50	200
test	50	50	50	50	200
total	495	924	391	348	2158

Tab. 1: Number of sources per class in the data sets.

## 3 Wasserstein GAN

Generative adversarial networks (GANs) [Go14; Sa16] are able to learn a representation of the underlying statistical distributions of sets of images. Sampling from those representations may provide additional data points for subsequent treatments [Bu21a; Bu22].

For this project, we employ a variant of the standard GAN setup called Wasserstein GAN (wGAN), which uses the Wasserstein-1 metric, often referred to as the Earth Mover’s distance, as main term in the loss function [ACB17]. A direct advantage of this setup is the correlation between image quality and the value of the loss function, transforming the discriminant of a standard GAN into a critic. Additionally, training of wGANs is often more stable and more likely to converge than standard GAN setups. To approximate the Wasserstein-1 metric by use of a critic network, it has to be ensured that the 1-Lipschitz constraint is fulfilled. This is achieved by applying a gradient penalty term to the loss function as in [Gu17].

Since different morphologies of radio galaxies result in very different images, it seems reasonable to condition the networks with the class label. For this setup, this is achieved for the generator by applying a 2D transposed convolution operator on a matrix of image dimensions filled with the class label. The transpose-convoluted layer is then concatenated to the first transpose-convoluted layer of the noise tensor. Batch normalisation in 2D and ReLU activation functions are used. The concatenated tensor is then passed through five additional 2D transposed convolutions, where no normalisation or activation is applied

after the last layer. Instead, the individual pixel values are simply clipped to  $[-1,1]$  for easy conversion to greyscale. The critic is built analogously, but uses 2D convolutional layers resulting in a single output node representing the critic score for image quality. Here, layer normalisation and leaky ReLU functions are used except for the last layer. A schematic of the wGAN setup can be found in Fig. 2.

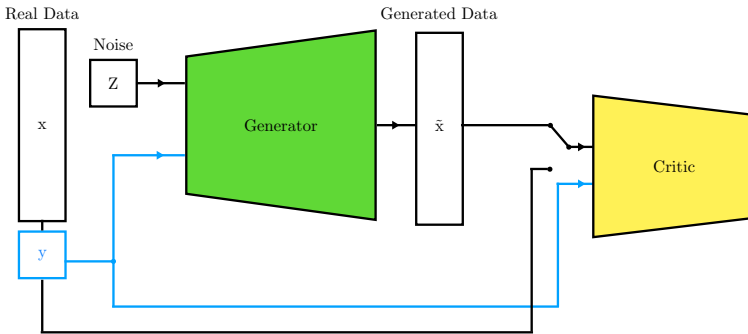


Fig. 2: Schematic of the wGAN architecture, where  $x$  denotes real images and  $y$  the class label. The noise vector  $Z$  is the input of the generator and the output of the generator  $\tilde{x}$  represents generated images. The critic receives either real or generated images as input.

## 4 Generated Images

### 4.1 Training

To evaluate statistical fluctuations, ten statistically independent wGAN training runs are launched solely on the training set. The training is performed with a single NVIDIA A100 GPU provided by the Maxwell cluster at DESY. The training is done for 40,000 generator iterations, while the critic is trained five times per generator iteration. A batch size of 400 is chosen and the full training takes roughly 7 hours to complete. The generator and critic weights are saved every 250 iterations, allowing to scan for the best training state later on.

### 4.2 Image Quality

In order to determine the quality of images and thus to find the best performing training iteration, a set of distributions is defined to compare generated images to the validation data set. This includes normalised histograms of pixel intensities, the number of pixels with an intensity greater than zero and of the sum of intensities. These histograms are compared for each class and the relative mean absolute error (RMAE) between the generated set of 10,000 images and the validation set is computed. The RMAE for the different histograms are

summed up to yield a single figure-of-merit (FOM) per class, where the training iteration with the lowest FOM value is used in the following. As an example, the histogram of pixel intensities for FRI is depicted in Fig. 3, where the generated histogram is shown in blue and the validation set in orange.




Fig. 3: Histogram of pixel intensities of FRI sources comparing batches of real (orange) and generated images (blue) for the training epoch with the lowest combined RMAE. The per-bin relative error is shown in the bottom panel.

For a direct visual comparison of image quality, a set of 5,000 images per class is generated and aligned using principle component analysis. Subsequently, the pixel-by-pixel difference is computed for all possible pairs of real and generated images. The closest pairs with a minimal activity for each class are shown in Fig. 4.

## 5 Using wGAN-supported Augmentation

As described above, classical augmentation relies on rotation and reflection of training images, adding no statistical information on the morphology. This type of augmentation is, thus, mainly a way to help a classifier understand the rotational invariant nature of radio galaxy images. Nowadays, this could be achieved by directly using equivariant models that incorporate such symmetry constraints inherently and show potential to improve classifier training without relying exclusively on augmentation [Bo21].

Here, we evaluate the augmentation achieved by using a combined data set of real and generated images (i.e. wGAN-supported augmentation) by comparing it to using the classically augmented (real-only) data set. The combined data set is created before starting the classifier training with a fixed ratio between the number of generated  $i_g$  and real images  $i_r$ , denoted  $\lambda = i_g/i_r$ . We study the range  $\lambda = 1, \dots, 4$  and the class distribution of the




Fig. 4: Closest matching pairs of generated and real images.

generated images is designed such that the combined data set is balanced. For the real-only training runs, class weights are introduced to the loss function to account for the class imbalance in the data set.

We train a simple classifier, namely a fully-connected neural network with three hidden layers and four output nodes. The classifier is trained for 12,500 epochs with a batch size of 100 and a learning rate of  $5 \cdot 10^{-4}$ . Training takes on average 3.5 seconds per epoch on an NVIDIA V100 GPU for real-only runs. This setup is not optimised to reach maximal classification accuracies as the goal of this study is only to compare classical with wGAN-supported augmentation.

We observe a significant improvement when increasing the number of generated images in the combined data set. As shown in Fig. 5 the real-only training run enters into over-training (i.e. the value of the loss function increases, while the training accuracy deviates further from the validation accuracy) much earlier than training runs using a combined training data set. The start of the over-trained period depends on  $\lambda$ , where we observe that adding more generated images to the data set shifts the start to later training iterations, while the validation loss score continues to decrease and the validation accuracy score (not shown) increases.

In order to compare the overall performance between different training setups, the multi-class Brier score [Br50] is used to determine the best training iteration for each of the ten statistically independent training runs for our five different training setups. The Brier score is essentially the mean squared error of the predicted probabilities of a classifier for all classes. This has the advantage that also the certainty of the classifier's decision is considered, which winner-takes-all FOMs such as accuracy do not take into account. The performance of the best models (i.e. models with minimal Brier score) is evaluated on

an independent test set that contains real data only using commonly applied metrics such as accuracy, precision, recall, and  $F_1$  score. The  $F_1$  score is the harmonic mean between precision and recall. The results for the  $F_1$  score are depicted in Fig. 6, indicating that training runs using the combined data set with  $\lambda = 4$  achieve an  $F_1$  score that is  $(23 \pm 2) \%$  higher than for the real-only training runs.




Fig. 5: Validation loss during training of the classifier for different values of  $\lambda$ . The solid line represents the mean of ten independent training runs and the shaded area represents the symmetrised standard deviation.




Fig. 6: Performance of the best model on an independent test set. We show the mean of ten independent training runs, where the uncertainty is given as the symmetrised standard deviation.

## 6 Discussion

Previous approaches employing generative models to simulate images of radio galaxies are based on variational autoencoders (VAE) [Ba21; Ma18; Ma19]. Therefore our GAN-based approach is novel in the field of radio astronomy. In this work, we demonstrate that we

are able to generate highly realistic, artificial images of radio sources. The quality of the generated images is evaluated in different ways and with different metrics. After visual inspection of a large number of generated images, our generated images do not suffer from issues known from other state-of-the-art simulations based on generative networks. In summary, we find:

- The resolution of the generated images is as good as the resolution of the images we train on.
- The generated images have a similar noise level as the preprocessed training examples.
- We do not observe any pseudo-textures or pseudo-structures in the generated images.

The images generated with the wGAN can be put to use for applications such as classifier training. As a result, our contribution represents a major improvement for simulations of radio galaxies based on generative networks. Butter et al. [Bu21b] showed for toy models that an amplification of training statistics with generative models is possible. Our work represents an example of this GANplyfication with real astrophysical data. For a recent particle physics application see [Bi22].

The presented method allows to amplify the performance of a simple classification model substantially by an increase in the  $F_1$  score of 23 % for a ratio of  $\lambda = 4$  between generated and real images and can be particularly useful for applications with unbalanced data sets. The application of wGAN-supported augmentation to more complex classifiers such as CNNs is left for future work.

## Acknowledgements

This work was supported by UHH, DESY, TUHH and HamburgX grant LFF-HHX-03 to the Center for Data and Computing in Natural Sciences (CDCS) from the Hamburg Ministry of Science, Research, Equalities and Districts. This project benefits greatly from the exchange with particle physicists with a vast experience in using generative models for calorimeter simulations and was supported in part through the Maxwell computational resources operated at DESY.

## References

- [ACB17] Arjovsky, M.; Chintala, S.; Bottou, L.: Wasserstein GAN, 2017, arXiv: 1701.07875.
- [AT17] Aniyani, A. K.; Thorat, K.: Classifying Radio Galaxies with the Convolutional Neural Network. The Astrophysical Journal Supplement Series 230/2, p. 20, June 2017.



- 
- [ATV18] Alhassan, W.; Taylor, A. R.; Vaccari, M.: The FIRST Classifier: compact and extended radio galaxy classification using deep Convolutional Neural Networks. *MNRAS* 480/2, pp. 2085–2093, July 2018.
  - [Ba21] Bastien, D. J.; Scaife, A. M. M.; Tang, H.; Bowles, M.; Porter, F.: Structured variational inference for simulating populations of radio galaxies. *MNRAS* 503/3, pp. 3351–3370, Mar. 2021.
  - [BCM17] Baldi, R. D.; Capetti, A.; Massaro, F.: FR0CAT: a FIRST catalog of FR 0 radio galaxies. *A&A* 609/, A1, Dec. 2017.
  - [Bi22] Bieringer, S.; Butter, A.; Diefenbacher, S.; Eren, E.; Gaede, F.; Hundhausen, D.; Kasieczka, G.; Nachman, B.; Plehn, T.; Trabs, M.: Calomplification – The Power of Generative Calorimeter Models, 2022, arXiv: 2202.07352.
  - [Bo21] Bowles, M.; Bromley, M.; Allen, M.; Scaife, A.: E(2) Equivariant Self-Attention for Radio Astronomy, 2021, arXiv: 2111.04742.
  - [Br50] Brier, G. W.: Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78/1, pp. 1–3, 1950.
  - [Bu21a] Buhmann, E.; Diefenbacher, S.; Eren, E.; Gaede, F.; Kasieczka, G.; Korol, A.; Krüger, K.: Getting High: High Fidelity Simulation of High Granularity Calorimeters with High Speed. *Comput Softw Big Sci* 5/13, May 2021.
  - [Bu21b] Butter, A.; Diefenbacher, S.; Kasieczka, G.; Nachman, B.; Plehn, T.: GANplying event samples. *SciPost Physics* 10/6, June 2021.
  - [Bu22] Buhmann, E.; Diefenbacher, S.; Eren, E.; Gaede, F.; Hundhausen, D.; Kasieczka, G.; Korcari, W.; Krüger, K.; McKeown, P.; Rustige, L.: Hadrons, Better, Faster, Stronger. *Mach. Learn.: Sci. Technol.* 3/2, July 2022.
  - [BWH95] Becker, R. H.; White, R. L.; Helfand, D. J.: The FIRST Survey: Faint Images of the Radio Sky at Twenty Centimeters. *The Astrophysical Journal* 450/, p. 559, Sept. 1995.
  - [Ca04] Carilli, C.; Furlanetto, S.; Briggs, F.; Jarvis, M.; Rawlings, S.; Falcke, H.: Probing the dark ages with the Square Kilometer Array. *New Astronomy Reviews* 48/11-12, pp. 1029–1038, Dec. 2004.
  - [CMB17a] Capetti, A.; Massaro, F.; Baldi, R. D.: FRICAT: A FIRST catalog of FR I radio galaxies. *A&A* 598/, A49, Jan. 2017.
  - [CMB17b] Capetti, A.; Massaro, F.; Baldi, R. D.: FRIICAT: A FIRST catalog of FR II radio galaxies. *A&A* 601/, A81, May 2017.
  - [Fr18] Frid-Adar, M.; Klang, E.; Amitai, M.; Goldberger, J.; Greenspan, H.: Synthetic Data Augmentation using GAN for Improved Liver Lesion Classification, 2018, arXiv: 1801.02385.
  - [FR74] Fanaroff, B. L.; Riley, J. M.: The morphology of extragalactic radio sources of high and low luminosity. *MNRAS* 167/, 31P–36P, May 1974.

- [GBW10] Gendre, M. A.; Best, P. N.; Wall, J. V.: The Combined NVSS-FIRST Galaxies (CoNFIG) sample - II. Comparison of space densities in the Fanaroff-Riley dichotomy. *MNRAS* 404/4, pp. 1719–1732, Mar. 2010.
- [Go14] Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y.: Generative Adversarial Networks, 2014, arXiv: 1406.2661.
- [Gu17] Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A.: Improved Training of Wasserstein GANs, 2017, arXiv: 1704.00028.
- [GW08] Gendre, M. A.; Wall, J. V.: The Combined NVSS-FIRST Galaxies (CoNFIG) sample - I. Sample definition, classification and evolution. *MNRAS* 390/2, pp. 819–828, Sept. 2008.
- [JM16] Jonas, J.; MeerKAT Team: The MeerKAT Radio Telescope. In: *MeerKAT Science: On the Pathway to the SKA*. 1, p. 1, Jan. 2016.
- [Ma18] Ma, Z.; Zhu, J.; Li, W.; Xu, H.: Radio Galaxy Morphology Generation using Residual Convolutional Autoencoder and Gaussian Mixture Models. In: 2018 25th IEEE ICIP. Pp. 3044–3048, 2018.
- [Ma19] Ma, Z.; Zhu, J.; Zhu, Y.; Xu, H.: Radio Galaxy Morphology Simulation via Residual Conditional Variational Autoencoder. In: 2019 15th CIS). Pp. 151–155, 2019.
- [MB17] Miraghaei, H.; Best, P. N.: The nuclear properties and extended morphologies of powerful radio galaxies: the roles of host galaxy and environment. *MNRAS* 466/4, pp. 4346–4363, Jan. 2017.
- [Pr11] Proctor, D. D.: Morphological Annotations For Groups In The FIRST Database. *The Astrophysical Journal Supplement Series* 194/2, p. 31, May 2011.
- [Sa16] Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X.: Improved Techniques for Training GANs, 2016, arXiv: 1606.03498.
- [Sa21] Samudre, A.; George, L. T.; Bansal, M.; Wadadekar, Y.: Data-Efficient Classification of Radio Galaxies. *MNRAS* 509/2, pp. 2269–2280, Oct. 2021.
- [TSL19] Tang, H.; Scaife, A. M. M.; Leahy, J. P.: Transfer learning for radio galaxy classification. *MNRAS* 488/3, pp. 3358–3375, Sept. 2019.
- [va13] van Haarlem, M. P. et al.: LOFAR: The LOw-Frequency ARray. *A&A* 556/A2, Aug. 2013.
- [Zh17] Zhu, X.; Liu, Y.; Qin, Z.; Li, J.: Data Augmentation in Emotion Classification Using Generative Adversarial Networks, 2017, arXiv: 1711.00648.