

Anomaly Detection in Radio Galaxy Data With Trainable Combination of Shifted Filter Response (COSFIRE) Filters

Steven Ndung'u, Trienko Grobler, Stefan J. Wijnholds, and George Azzopardi

Abstract – Detecting anomalies in radio astronomy is challenging due to the vast amounts of data and the rarity of labeled anomalous examples. Addressing this challenge requires efficient methods capable of identifying unusual radio galaxy morphologies without relying on extensive supervision. This work introduces an innovative approach to anomaly detection on the basis of morphological characteristics of the radio sources by using trainable Combination of Shifted Filter Response (COSFIRE) filters as an efficient alternative to complex deep learning methods. The framework integrates COSFIRE descriptors with an unsupervised local outlier factor (LOF) algorithm to identify unusual radio galaxy morphologies. Evaluations on a radio galaxy benchmark data set demonstrate strong performance, with the COSFIRE-based approach achieving a geometric mean score of 79%, surpassing the 77% achieved by a computationally intensive deep learning autoencoder. By characterizing normal patterns and detecting deviations, this semisupervised methodology overcomes the need for anomalous examples in the training set, a major limitation of traditional supervised methods. This approach shows promise for next-generation radio telescopes, where fast processing and the ability to discover unknown phenomena are crucial.

1. Introduction

The advent of next-generation radio telescopes has ushered in an era of unprecedented data volume in radio astronomy, with facilities such as the Square Kilometre Array (SKA) expected to annually generate exabytes of data [1]. Although processing such vast data

sets poses significant challenges, it also presents unique opportunities for serendipitous discoveries of novel astrophysical phenomena. The systematic exploration of these data sets necessitates robust anomaly detection methodologies capable of identifying both rare, theoretically anticipated phenomena and entirely novel, unforeseen ones.

Machine learning approaches, particularly those focused on learning compact representations of typical radio sources, have emerged as promising tools for this task. Autoencoders (AEs) have been successfully applied to anomaly detection across diverse astronomical data sets [2, 3], but the high computational cost and opaque latent representations present significant drawbacks. To address these challenges, we use trainable Combination of Shifted Filter Response (COSFIRE) filters [4], which are efficient, explainable and rotation tolerant, and whose responses form the basis of a feature descriptor [5]. COSFIRE filters are automatically configured to detect and describe the spatial arrangement of emission blobs, cores, jets, and lobes, which characterize a radio source's morphology. This approach has shown notable efficacy, accuracy, and robustness in radio galaxy classification and image retrieval [5, 6], outperforming deep learning approaches in certain cases. Notably, COSFIRE's rotation invariance ensures that galaxies with the same intrinsic morphology receive consistent labels, regardless of the apparent orientation due to the telescope's line of sight. Figure 1 shows examples of typical radio galaxies: Fanaroff–Riley Class I (FRI); Fanaroff–Riley Class II (FR II) [7]; real anomalous sources (X-shaped radio galaxy [XRG], and ring radio galaxy [RRG]) [8]; and synthetic anomalous sources used in this work (see Section 2.1).

2. Data and Methodology

2.1 Data Set

We evaluate our anomaly detection approach by using the Faint Images of the Radio-Sky at Twenty Centimeters (FIRST) Radio Galaxy Anomaly Detection Benchmark (FRGADB) data [9, 10]. FRGADB is composed of ~800 samples with four radio galaxy classes: FRI, FR II, XRG, and RRG, distributed between training, validation, and test sets (Table 1) [7, 8]. The images are 150 pixels \times 150 pixels, with radio sources centered. The images were preprocessed by using a morphologically augmented sigma clipping algorithm, which effectively

Part of this work is supported by the Foundation for Dutch Scientific Research Institutes. This work is based on the research supported, in part, by the National Research Foundation of South Africa (grants 119488 and CSRP2204224243). The financial assistance of the South African Radio Astronomy Observatory (SARAO) toward this research is hereby acknowledged (www.ska.ac.za).

Manuscript received 19 February 2025.

Steven Ndung'u and Trienko Grobler are with Computer Science, University of Stellenbosch, Cnr Banhoek Road & Joubert Street, Stellenbosch 7600, South Africa; e-mail: 26846578@sun.ac.za. e-mail: tlgrobler@sun.ac.za.

Stefan J. Wijnholds is with ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD. Dwingeloo, The Netherlands; e-mail: wijnholds@astron.nl.

George Azzopardi is with the Bernoulli Institute for Mathematics, Computer Science, and Artificial Intelligence, University of Groningen, Nijenborgh 9, 9747AG Groningen, The Netherlands; e-mail: url.azzopardi@rug.nl.

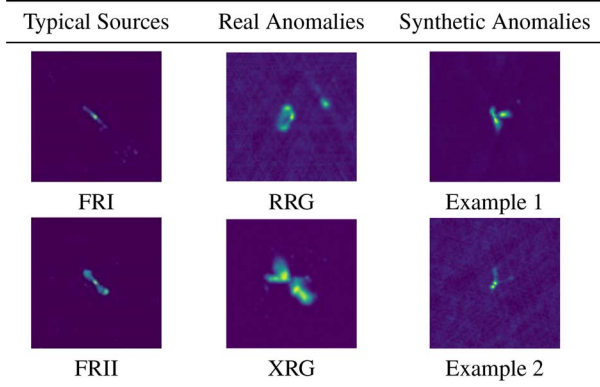


Figure 1. Examples of typical radio galaxies and both real and synthetic anomalous radio sources: FRI (edge-darkened source); FRII (edge-brightened source); XRG (x-shaped source); RRG (ring-shaped source); and synthetic sources with nonstandard morphologies.

filtered out noise, while preserving astronomical source pixels [3].

The presence of XRG and RRG anomalies in both the validation and test sets introduces potential data leakage, meaning the model is exposed during validation to patterns that closely resemble those it will encounter during testing, which can lead to an overestimation of its true generalization performance. To address this, we generated 90 synthetic anomalous samples from training images to replace the 22 XRG and 12 RRG galaxies in the given validation set [11]. The new validation set includes 22 FRI, 53 FRII, and 90 anomalous patterns. The 90 synthetic sources were generated through a systematic process using the training set. First, 10 FRI image pairs were randomly selected without replacement. One image from each pair underwent rotations of 90° , 180° , and 270° , followed by superimposition with its partner by using maximum pixel values. This created 30 anomalous sources. The same procedure was applied to 10 FRII pairs and 10 hybrid pairs (FRI-FRII combinations), yielding a total of 90 anomalous images (see examples in the third column of Figure 1). This procedure ensures that no prior knowledge of the test anomalies is leaked during training or validation, leading to a more reliable assessment of model performance.

2.2 Method

In this section, we sketch the configuration of COSFIRE filters and the derived feature vectors. A more detailed description can be found in [5]. The responses from a set of COSFIRE filters are used in a descriptor-based approach to characterize radio sources by morphological features. This approach is based on the analysis of the shape properties and is automatically configured by applying a series of center-on and center-off difference of Gaussian (DoG) filters to an image. A DoG function approximates the second-order derivative of a Gaussian (Laplacian of Gaussian), offering the

Table 1. Distribution of the FRGADB data set across training, validation, and test sets for each galaxy type (note that XRG and RRG are used only in the test set)

Galaxy type	Total samples	Training	Validation	Test
Regular classes				
FRI	211	137	22	22
FRII	523	387	53	53
Synthetic anomalies				
Synthetic			90	
Real anomalous classes (test only)				
XRG	44			22
RRG	25			13
Total	803	524	165	110

advantage of efficient computation due to its separability. These DoG filters come in two variants: center-on filters that detect transitions from bright to dark regions and center-off filters that identify dark to bright transitions. These filters, varying in scale and polarity (i.e., center-on or center-off), are applied to the input image to produce DoG response maps through convolution operations.

To configure a COSFIRE filter, salient features, such as radio source blobs, jets, or lobes, are detected along concentric circles on a given example training image (e.g., a specific radio source type). Key points are identified by locating the local maxima of the DoG responses along these circles. The number of circles and the respective radii act as hyperparameters of the filter. For each detected key point i , the following attributes are recorded: distance (ρ_i) and angle (ϕ_i) from the pattern center; filter polarity (δ_i); and the outer Gaussian's standard deviation (σ_i). These parameters are compiled into a set of tuples, denoted by C_f :

$$C_f = \{(\rho_i, \phi_i, \sigma_i, \delta_i) \mid i = 1, \dots, n\} \quad (1)$$

where n represents the total detected key points. In our data set, radio sources are already centered in the 150×150 images, so we use the image center as the COSFIRE reference point. In principle, one could instead select any point, such as the source's brightness centroid.

For anomaly detection, COSFIRE filters are first automatically configured to be selective for some randomly chosen sources drawn from the training set (refer to [5] and [6] for some examples). The configuration procedure defines the selectivity of a COSFIRE filter on the basis of the spatial relationships among the selected key points. These configured filters are expected to produce strong responses in images with radio structures similar to those seen during training and weak responses to unfamiliar patterns, thereby making them suitable to detect anomalies.

The application of a configured COSFIRE filter follows three steps. First, DoG response maps are computed for each unique combination of polarity (δ) and scale (σ).

Table 2. Hyperparameter search space in COSFIRE filter optimization; this table outlines the parameter ranges explored to identify optimal filter configurations^a

Parameter name	Values
σ : DoG outer standard deviation	{5, 6}
P : set of concentric circle radii	$\left\{ \{0, 5, 10, 15, 20, 25\}, \{0, 5, 10, 15, 20, 25, 30\} \right\}$
t_1 : response threshold	{0.05, 0.1}
σ_0, α : parameters for positional tolerance	$\sigma_0 \in \{0.5, 0.75\}$ $\alpha \in \{0.1, 0.15\}$

^a The positional tolerance $\hat{\sigma}_i$ for each filter subunit is computed as $\hat{\sigma}_i = \sigma_0 + \alpha\rho_i$, where ρ_i is the radial distance of the subunit from the filter center. This value $\hat{\sigma}_i$ is used as the standard deviation of a Gaussian function that controls the spatial tolerance during filter response computation.

Second, these response maps are blurred with a Gaussian function to account for positional variations. See Table 2 for the parameters used to compute the standard deviation $\hat{\sigma}_i$ of the Gaussian, which increases linearly with the distance ρ_i from the filter center. This function introduces spatial tolerance in the filter response. For more details, please refer to [5]. Without this blurring step, the COSFIRE filter would only respond to exact replicas of the pattern in the training image used for the filter configuration. Therefore, the blurring introduces some tolerance to spatial deviations, enabling the filter to recognize patterns that slightly differ from the preferred configuration. Each blurred response map is then shifted in the opposite direction of its corresponding polar coordinates (ρ_i, ϕ_i) so that all DoG responses meet at the filter center [5]. Finally, the COSFIRE filter response is obtained by calculating the geometric mean (G-Mean) across all aligned and blurred response maps. Rotational invariance is achieved by generating multiple rotated versions of the base COSFIRE filter, each created by offsetting the angular parameter ϕ_i of every key point by an angle $\psi \in \{0, \frac{\pi}{6}, \dots, 11\frac{\pi}{6}\}$. A rotation-invariant COSFIRE response is then obtained by taking the maximum response across all these rotated filters. The final COSFIRE descriptor, a numerical feature vector representing the structural patterns in an image, is constructed by applying all rotation-invariant filters to an image and recording the maximum responses. This

feature vector effectively captures the essential structural elements of a given radio galaxy image. Following [5], the descriptors are then L2 normalized.

We then used the local outlier factor (LOF) algorithm [12] to detect anomalies in COSFIRE descriptors of radio galaxies. The full pipeline is shown in Figure 2.

2.3 Evaluation Metrics

To evaluate our anomaly detection framework, we use metrics specifically designed for imbalanced data sets, as anomalous radio galaxy sources represent a small fraction of the data set. We assess model performance by using four key metrics: *precision*, *recall*, *specificity*, and *G-Mean*. Precision quantifies the proportion of true anomalies among all instances predicted as anomalies. Recall measures the model’s ability to correctly identify actual anomalous sources. Specificity evaluates the model’s accuracy in recognizing normal (non-anomalous) galaxies. G-Mean is calculated as the square root of the product of specificity and recall, and it reflects the model’s balanced performance across both majority and minority classes. These metrics collectively offer a comprehensive assessment of the model’s ability to distinguish rare, anomalous radio sources, while maintaining reliable detection of typical galaxy morphologies. Note that the performance evaluation considers anomalous sources as positive instances and normal sources as negative instances in all metric calculations.

3. Experimental Results

3.1 Performance

To evaluate the performance of our COSFIRE filter approach in anomaly detection, we adopted an experimental design closely aligned with the methodologies presented in [5] and [6]. To ensure representative sampling across classes, the training data was split by using a stratified holdout approach, randomly dividing the data set into training and validation sets, while maintaining the original class proportions. Prior work

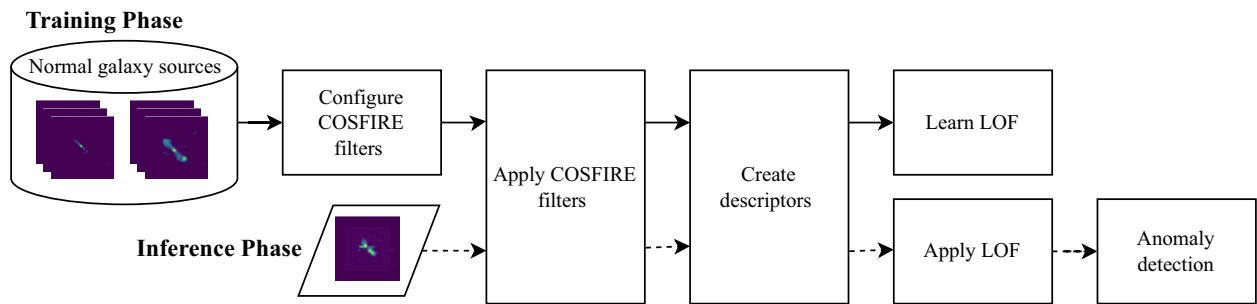


Figure 2. A schematic illustration of the proposed COSFIRE framework for detecting anomalous radio galaxies, showing training and inference phases of the workflow. The training phase includes three key steps: configuring the COSFIRE filters, extracting descriptors from known radio galaxy images, and training a LOF model for anomaly detection. During the inference phase, new radio images undergo the same descriptor extraction before being evaluated by the trained LOF model to identify potential anomalies.

[5, 6] demonstrated that COSFIRE filter-generated descriptors retain robustness across varied hyperparameter spaces. On the basis of these findings, we simplified our workflow by reducing the hyperparameter search space, selecting two values per COSFIRE hyperparameter, instead of three, as done in [5] (Table 2).

For our unsupervised learning approach, we optimized the LOF algorithm by systematically exploring its hyperparameters: leaf size $\in \{40, 60, 80\}$, contamination $\in \{0.37, 0.38, 0.39\}$, number of neighbors $\in \{8, 9, 10\}$, and metric $\in \{\text{cosine, Manhattan, Minkowski}\}$ [12]. All remaining parameters were maintained at the default values to constrain the optimization space. Therefore, each COSFIRE configuration underwent 81 distinct experiments during the hyperparameter grid search. We performed a full grid search over LOF hyperparameters for each of the 32 COSFIRE configurations, yielding an 81×32 matrix of evaluation scores.

We first selected COSFIRE configurations by computing mean G-Mean scores across the validation set. A right-tailed Student's t -test was used to compare the highest performing descriptor against all other configurations [13]. Descriptors achieving P values ≥ 0.05 were considered statistically similar to the optimal configuration and retained for further analysis. This filtering identified two descriptors that performed consistently well, reducing the evaluation matrix to 81×2 . These descriptors share the same parameter values $\sigma = 6$, $P = \{0, 5, 10, 15, 20, 25, 30\}$, $\sigma'_0 = 0.75$, and $\alpha = 0.15$ and differ only in the value of t_1 (0.10 and 0.05). For each of these descriptors, the top 10 performing LOF hyperparameter sets were selected and evaluated on the test set.

This analysis considered varying numbers of COSFIRE filters, ranging from 5 to 100 in the validation set. Because the configuration was performed on two distinct known classes, FRI and FRII (refer to Section 2.2), the dimensionality of the feature descriptor proportionally scales with the number of filters, where n filters produced a $2n$ -dimensional descriptor (e.g., five filters per class generated a 10-dimensional descriptor, six filters a 12-dimensional descriptor, and so forth). The experiments revealed that 90 COSFIRE filters per class were required to achieve optimal G-Mean performance on the validation set.

3.2 Comparison With Previous Work

The results of this study are compared with the state-of-the-art findings presented in [3], which used the same data set. This work used principal component analysis (PCA) for feature extraction combined with the LOF for anomaly detection. Also, they investigated three convolutional autoencoder (CAE) architectures: standard CAE (SCAE), baseline CAE (BCAE) [2], and memory unit SCAE (MemSCAE); see [3, 14]. BC AE serves as a shallow network autoencoder that reconstructs input data by using reconstruction error as the

Table 3. Performance comparison of various models using LOF for anomaly detection in a semisupervised setting

Model	G-Mean	Specificity	Recall
PCA	0.60 ± 0.05	0.78 ± 0.03	0.47 ± 0.07
MemSCAE	0.77 ± 0.02	0.68 ± 0.07	0.90 ± 0.06
SCAE	0.77 ± 0.01	0.69 ± 0.05	0.85 ± 0.05
BCAE	0.72 ± 0.02	0.72 ± 0.09	0.73 ± 0.08
COSFIRE	0.79 ± 0.01	0.70 ± 0.00	0.89 ± 0.01
COSFIRE ^a	0.78 ± 0.03	0.70 ± 0.02	0.88 ± 0.04

^a The COSFIRE model used the synthetic validation set for hyperparameter tuning.

loss function. SCAE enhances upon BC AE by applying a deeper network [3]. MemSCAE further extends the architecture with a memory unit that constrains the latent space to some degree, thereby improving anomaly detection capabilities [14].

The extensive experiments conducted in [3] demonstrated that MemSCAE and BC AE achieved a consistent G-Mean performance of ~ 0.77 in anomaly detection tasks, outperforming both PCA + LOF and SCAE approaches. Our COSFIRE + LOF framework surpassed these benchmarks, achieving an average G-Mean of ~ 0.79 on the test data set (Table 3). In addition, when using the synthetic validation set (see Section 2.1) for hyperparameter tuning, we obtained a G-Mean score of 0.78 ± 0.03 by using 67 COSFIRE filters per class.

4. Discussion and Conclusion

The COSFIRE-based methodology generates robust feature representations that effectively detect anomalies not present in the training data, while maintaining high performance despite the challenges of limited training data and class imbalance. This results in improved performance, with a G-Mean score of 79%, surpassing both the competing deep learning autoencoder (77%) and the PCA-based approach (60%) reported in [3], and with the latter being a similarly lightweight anomaly detection method.

By leveraging the COSFIRE methodology, our approach addresses the scalability challenge of anomaly detection in next-generation radio surveys by using lightweight, rotation-invariant descriptors that outperform deep autoencoder pipelines on the FIRST benchmark data set. We refer readers to our prior work [5, 6] for detailed floating point operation comparisons demonstrating COSFIRE's efficiency versus deep learning methods. Its demonstrated effectiveness highlights particular promise for enabling serendipitous discoveries with advanced telescopes such as SKA, where efficient processing and the ability to uncover unknown phenomena are crucial. However, the evaluation is limited by the modest size of the FRGADB data set (803 images) and its focus on only two anomalous classes (XRG and RRG). Future work will expand validation to larger and more diverse catalogs and additional rare morphologies.

5. Acknowledgments

We thank the Center for Information Technology of the University of Groningen for support and for providing access to the Hábrók high-performance computing cluster. The code and data set used in this work are publicly available on GitHub at https://github.com/stevenndungu/anomaly_detection. A brief preview of the analyses can be found at https://stevenndungu.github.io/anomaly_detection/.

6. References

1. J. Farnes, B. Mort, F. Dulwich, S. Salvini, and W. Armour, "Science Pipelines for the Square Kilometre Array," *Galaxies*, **6**, 4, November 2018, p. 120.
2. F. Ventura, *Prospecting for Enigmatic Radio Sources With Autoencoders: A Novel Approach*, Master's thesis, University of Pretoria, South Africa, 2022.
3. K. Brand, T. L. Grobler, and W. Kleynhans, "CARA: Convolutional Autoencoders for the Detection of Radio Anomalies," *RAS Techniques and Instruments*, **4**, February 2025, p. rzaf005.
4. G. Azzopardi and N. Petkov, "Trainable COSFIRE Filters for Keypoint Detection and Pattern Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **35**, 2, February 2012, pp. 490-503.
5. S. Ndung'u, T. Grobler, S. J. Wijnholds, D. Karastoyanova, and G. Azzopardi, "Classification of Radio Galaxies with Trainable COSFIRE Filters," *Monthly Notices of the Royal Astronomical Society*, **530**, 1, May 2024, pp. 783-794.
6. S. Ndung'u, T. Grobler, S. J. Wijnholds, and G. Azzopardi, "Content-Based Image Retrieval Using COSFIRE Descriptors With Application to Radio Astronomy," *Monthly Notices of the Royal Astronomical Society*, **537**, 4, March 2025, pp. 3286-3297.
7. 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*, **167**, 1, April 1974, pp. 31P-36P.
8. D. D. Proctor, "Morphological Annotations for Groups in the FIRST Database," *The Astrophysical Journal*, **194**, 2, May 2011, p. 31.
9. R. H. Becker, R. L. White, and D. J. Helfand, "The FIRST Survey: Faint Images of the Radio Sky at Twenty Centimeters," *The Astrophysical Journal*, **450**, September 1995, p. 559.
10. K. Brand, *FRGADB—FIRST Radio Galaxy Anomaly Detection Benchmark*, version 1.0, <https://doi.org/10.5281/zenodo.13773680>.
11. H. M. Schlüter, J. Tan, B. Hou, and B. Kainz, "Natural Synthetic Anomalies for Self-Supervised Anomaly Detection and Localization," European Conference on Computer Vision, **13691**, October 2022, pp. 474-489.
12. M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying Density-Based Local Outliers," ACM Special Interest Group on Management of Data International Conference on Management of Data, Dallas, TX, USA, May 16-18, 2000, pp. 93-104.
13. G. Marshall and L. Jonker, "An Introduction to Inferential Statistics: A Review and Practical Guide," *Radiography*, **17**, 1, February 2011, pp. e1-e6.
14. D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, et al., "Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection," In Proceedings of the IEEE/CVF international conference on computer vision, Seoul, Korea (South), October 27-29, 2019, pp. 1705-1714.