Interpretable Gravitation-Based Outlier Detection via Density-Guided Mass Adaptation

Agnieszka Duraj[1][0000−0002−3047−6662]
, Piotr S. Szczepaniak[1][0000−0002−9973−0673], and Daria Rogowska[2]

[1] Institute of Information Technology, Lodz University of Technology,
al. Politechniki 8, 93-590 Łódź, Poland
[2] Accenture Poland, University Business Park,
Wólczańska 178, 90-530 Lodz, Poland
{agnieszka.duraj,piort.szczepaniak}@p.lodz.pl

Abstract. Centroid-based classifiers are attractive due to their interpretability and low computational cost, yet their performance degrades in the presence of class imbalance, irregular decision boundaries, and local density variations. This paper introduces DG-MA (Density-Guided Mass Adaptation), a gravitation-based model that incorporates local density information into the global structure of the Gravitation Model. The method applies a single batch adaptation of class masses derived from density-based structural statistics, avoiding iterative coupling between clustering and classification. As a result, DG-MA preserves the simplicity of gravitational decision rules while improving robustness to non-uniform data distributions in supervised outlier detection. Experimental results on multiple benchmark datasets show that DG-MA consistently improves over the classical Gravitation Model and remains competitive with density-based methods, while maintaining favourable computational efficiency.

Keywords: gravitation model · outlier detection · local density

1 Introduction

Outlier detection in high-dimensional and imbalanced data remains challenging due to local density variation, irregular class geometry, and the need for computationally efficient decision models. In this study, we consider supervised outlier detection formulated as a binary classification task, where normal observations form the negative class and outliers form the positive class.

Centroid-based models are commonly used due to their simplicity and low computational cost [1]. In particular, the Gravitation Model (GM) represents each class by a centroid and a mass parameter controlling class influence [2–4].

To address this limitation, we propose Density-Guided Mass Adaptation (DG-MA), a hybrid extension of GM in which DBSCAN [5] is used only during training to extract density-based structural statistics. These statistics are then aggregated into a single batch correction of class masses, after which inference remains purely centroid-based. As a result, the method combines local

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

density information with a transparent gravitational decision rule while preserving linear-time prediction complexity. Unlike existing Gravitation Model extensions based on iterative learning or misclassification-driven updates, DG-MA introduces a non-iterative density-guided mass adaptation mechanism. This approach preserves analytical interpretability while incorporating local structural information into the model.

The main contributions of this paper are as follows:

- we introduce a density-guided extension of the Gravitation Model for supervised outlier detection;
- we show how density-induced structural errors translate into global mass-ratio adjustments and decision-boundary shifts;
- we experimentally demonstrate that DG-MA improves over GM on several benchmark datasets while maintaining low inference complexity.

The remainder of this paper is organised as follows. Section 2 reviews related work on outlier detection and gravity-based classification models. Section 3 introduces the proposed DG-MA framework and analyses its theoretical properties. The experimental setup and evaluation protocol are described in Section 4, while the obtained results are discussed in Section 5. Finally, Section 6 concludes the paper and outlines directions for future research.

2 Related Work

Outlier detection aims to identify observations that deviate from the underlying data distribution and correspond to rare events [6]. A wide range of approaches has been developed, differing in modelling assumptions, scalability, and interpretability [7, 8].

Statistical methods identify outliers as low-probability samples [9], but are sensitive to model mismatch and high dimensionality. Distance-based approaches (e.g., k-NN [10]) rely on neighbourhood dispersion, yet degrade in high-dimensional spaces.

Density-based methods such as LOF, COF, and DBSCAN [11, 12, 5] detect outliers in low-density regions and capture complex cluster shapes, but are parameter-sensitive and computationally demanding. Machine learning approaches, including one-class SVMs, autoencoders, and ensembles [13–15], model complex patterns but often lack interpretability.

Importantly, methods such as LOF, Isolation Forest, and one-class SVM operate at the instance level, focusing on local deviations rather than explicitly shaping global decision boundaries. Consequently, they do not address geometrically interpretable boundary recalibration in centroid-based models.

In contrast, gravity-based classifiers provide a global and interpretable framework. The Gravitation Model (GM) represents each class by a centroid and mass [2, 3], with extensions addressing imbalance and feature weighting [16, 17]. However, these approaches remain inherently global and rely on misclassification-driven updates.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

Relatively few methods combine global class representations with local density information while preserving interpretability and low inference cost. Existing hybrid approaches often rely on iterative clustering–classification coupling, increasing complexity.

The proposed DG-MA addresses this gap by performing a single density-guided mass correction using DBSCAN, followed by purely centroid-based inference. Unlike prior GM extensions, it uses density-induced structural discrepancies instead of classification errors, preserving interpretability and linear-time prediction.

3 Methods

3.1 Gravitation Model

The Gravitation Model (GM), proposed in [4], is a centroid-based classification method inspired by the law of gravitation. Each class C_j is represented by its centroid c_j and an associated mass M_j, which controls the influence of the class in the feature space. This formulation allows the model to account for class asymmetry and partially compensate for class imbalance.

For an object x, the attraction of class C_j is defined as

$$F(x, c_j) = M_j \cdot sim(x, c_j), \tag{1}$$

where $sim(\cdot, \cdot)$ denotes a similarity function. In this work, we use the regularised inverse Euclidean similarity

$$sim(x, c_j) = \frac{1}{\|x - c_j\|_2 + \varepsilon_{\text{sim}}}, \tag{2}$$

where $\varepsilon_{\text{sim}} > 0$ ensures numerical stability. The predicted label is assigned to the class with the largest attraction force.

Unlike the standard nearest-centroid rule, GM introduces class-specific masses, which provide a more flexible global decision mechanism.

3.2 Density-Guided Mass Adaptation (DG-MA)

DG-MA extends the Gravitation Model by incorporating local density information extracted using DBSCAN. Unlike classical GM learning, class masses are not updated iteratively but determined via a single density-guided correction step.

Let $D_{tr} \subset \mathbb{R}^d$ denote the training set with K classes. Each class C_j is represented by centroid c_j and mass M_j, and prediction follows the standard gravitational rule:

$$\hat{y}(x) = \arg \max_{1 \le j \le K} F(x, c_j). \tag{3}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

DBSCAN assigns cluster labels $L(x)$, with $L(x) = -1$ denoting noise. Each cluster ℓ is mapped to a class via majority voting:

$$\mathrm{map}(\ell) = \arg\max_j \left|\{x : \ L(x) = \ell \wedge y(x) = C_j\}\right|. \tag{4}$$

Two class-level structural errors are then defined:

$$\beta_{j1} = \left|\{x : \ y(x) = C_j, \ (L(x) = -1 \vee \mathrm{map}(L(x)) \neq C_j)\}\right|, \tag{5}$$

$$\beta_{j2} = \left|\{x : \ y(x) \neq C_j, \ L(x) \neq -1, \ \mathrm{map}(L(x)) = C_j\}\right|. \tag{6}$$

These quantify lack of support and contamination. A normalised error is given by

$$\tilde{\beta}_j = \frac{\beta_{j1} + \lambda\, \beta_{j2}}{|C_j|}, \tag{7}$$

and class masses are adapted as

$$M_j = \max\Big(M_{\min}, \ 1 - \alpha\, \tilde{\beta}_j\Big). \tag{8}$$

After this single update, DG-MA applies the standard GM decision rule, preserving interpretability and linear-time inference. Since DBSCAN is used only to extract aggregated statistics, the method is less sensitive to parameter variations than instance-level density models.

However, the hard majority mapping may over-penalise minority samples in heterogeneous clusters, potentially leading to over-correction.

3.3 Influence of the Mass Ratio on Decision Boundaries

For two classes C_p and C_q with centroids c_p, c_q and masses M_p, M_q, the decision boundary is defined by

$$M_p\, sim(x, c_p) = M_q\, sim(x, c_q). \tag{9}$$

Using the regularised inverse Euclidean similarity, this yields

$$\frac{\|x - c_p\|_2 + \varepsilon_{\mathrm{sim}}}{\|x - c_q\|_2 + \varepsilon_{\mathrm{sim}}} = \frac{M_p}{M_q}. \tag{10}$$

In the limit $\varepsilon_{\mathrm{sim}} \to 0^+$,

$$\frac{\|x - c_p\|_2}{\|x - c_q\|_2} = \frac{M_p}{M_q}, \tag{11}$$

which defines an Apollonius surface. Thus, the boundary geometry is fully determined by the mass ratio $\frac{M_p}{M_q}$.

In DG-MA, where $M_j = 1 - \alpha\tilde{\beta}_j$, density-induced structural errors directly control boundary shifts: higher $\tilde{\beta}_j$ reduces M_j and contracts the decision region, while lower values expand it.

Hence, DG-MA preserves the analytical form of the classifier while globally adjusting boundaries via mass adaptation.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

3.4 Computational Complexity Analysis

In GM, classification requires $O(K \cdot d)$ per instance, yielding $O(n \cdot d)$ for fixed K. DG-MA adds a single DBSCAN execution during training, with complexity ranging from $O(n \cdot d + n \log n)$ to $O(n \cdot d + n^2)$ [18]. Thus, training cost increases, but inference remains $O(K \cdot d)$. Although DG-MA reduces sensitivity to DBSCAN parameters through class-level aggregation, strongly suboptimal configurations may still affect mass estimation.

4 Experiments

The evaluation was conducted on five benchmark datasets reformulated as binary outlier detection tasks: Cardio [19], Vowels [20, 21], Musk [22], Chemical [23], and Synthetic [24]. Cardio contains 1831 instances with 21 features (9.6% outliers), Vowels 1456 instances with 12 features (3.4%), Musk 3062 instances with 166 features (3.2%), Chemical 3251 instances with 34 features (6.2%), and Synthetic 1000 instances with 10 features (5.0%).

We compared DBSCAN, the Gravitation Model (GM), and the proposed DG-MA. DBSCAN parameters (ε, $minPts$) were tuned on the training set via grid search maximising $F_{1,\text{out}}$, using an inner validation split within the training data. The same protocol was applied when DBSCAN was used within DG-MA.

All experiments used a 70:30 train–test split with 10 repeated stratified runs. Features were standardised using training-set statistics, and all parameter tuning, including DBSCAN parameters (ε, $minPts$) and DG-MA parameters (α, λ), was performed exclusively on the training data. Performance was evaluated using ROC-AUC, PR-AUC, $F_{1,\text{out}}$, TPR_{out}, and MCC, which together reflect ranking quality and rare-class detection under class imbalance. The same data splits were used for all compared methods to ensure a fair and consistent evaluation.

Although DBSCAN parameters depend on data density, in DG-MA they are used only to extract aggregated structural statistics from the training set. These statistics are computed at the class level and do not rely on precise cluster assignments in the test data. As a result, the method is less sensitive to density variations between training and test sets compared to standard density-based classifiers. Moreover, the use of repeated stratified splits reduces the impact of sampling variability on parameter estimation.

All evaluation metrics are maximised, i.e., values closer to 1 indicate better performance. In particular, $F_{1,\text{out}}$, PR-AUC, and ROC-AUC jointly reflect detection quality under class imbalance. All experiments were implemented in Python using scikit-learn, and code together with parameter settings will be made publicly available upon acceptance.

5 Results

All results are reported on the test set following the protocol described in Section 4. A comparison of DBSCAN, GM, and DG-MA is presented in Tab. 1.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

All evaluation metrics are maximised (values closer to 1 indicate better performance).

Table 1. Comparison of outlier detection performance for DBSCAN, GM, and DG-MA.

Metric	Dataset / Algorithm														
	Cardio			Vowels			Musk			Chemical			Synthetic		
	DBSCAN	GM	DG-MA	DBSCAN	GM	DG-MA	DBSCAN	GM	DG-MA	DBSCAN	GM	DG-MA	DBSCAN	GM	DG-MA
$F_{1,\text{out}}$	0.52	0.61	0.66	0.62	0.49	0.55	1.00	1.00	1.00	1.00	0.76	0.86	0.84	0.10	0.22
PR–AUC	0.56	0.63	0.64	0.64	0.50	0.55	1.00	1.00	1.00	1.00	0.78	0.88	0.86	0.16	0.38
ROC-AUC	0.71	0.77	0.78	0.77	0.72	0.74	1.00	1.00	1.00	1.00	0.78	0.88	0.88	0.47	0.59

As shown in Tab. 1, DBSCAN achieves near-perfect performance on Musk and Chemical (all metrics close to 1.0), but degrades under weak separability (Cardio), reflecting limitations of purely local density modelling. GM provides a stable and interpretable global model, performing well on separable datasets but degrading for heterogeneous structures (Vowels, Synthetic).

DG-MA consistently improves over GM in datasets with non-uniform density. The largest gains are observed for Cardio and Vowels, where higher $F_{1,\text{out}}$, PR-AUC, and ROC-AUC indicate improved rare-class detection. On Chemical, DG-MA significantly outperforms GM and approaches DBSCAN. For Musk, all methods achieve near-perfect results, while for Synthetic both GM and DG-MA remain limited.

Results correspond to mean values over 10 repeated splits, with standard deviations below 0.03 for all configurations. The analysis focuses on $F_{1,\text{out}}$, PR-AUC, and ROC-AUC as robust metrics under class imbalance. Additional comparison with classical methods is shown in Tab. 2.

On Cardio, One-Class SVM achieves the highest ranking performance, while DG-MA remains competitive with LOF. On Vowels, LOF dominates in ranking metrics, whereas DG-MA provides a more balanced precision–recall trade-off. Isolation Forest exhibits high recall but low precision. On Musk, DG-MA achieves near-perfect separation on this dataset. These results confirm that incorporating density information into centroid-based models improves robustness to heterogeneous data distributions while preserving interpretability.

DG-MA differs from instance-level methods by aggregating density information at the class level, enabling a transparent decision mechanism with linear-time inference. Its effectiveness depends on data structure: it is most beneficial for heterogeneous density and moderate overlap, while density-based methods remain preferable for well-separated data and performance remains limited for highly irregular distributions.

The additional computational cost arises from a single DBSCAN execution during training, whereas inference remains identical to GM ($O(K \cdot d)$). All baseline methods were evaluated using the same data splits and evaluation protocol for consistency.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

Table 2. Comparison of DG-MA with classical outlier detection methods.

Datasets	Algorithm	ROC-AUC	PR-AUC	$F_{1,\text{out}}$	TPR_{out}
Cardio	LOF	0.9437	0.6268	0.6000	0.6000
	Isolation Forest	0.8021	0.2110	0.1176	0.2000
	One-Class SVM	0.9803	0.8002	0.6804	0.9429
	DG-MA	0.78	0.64	0.66	0.65
Vowels	LOF	0.9897	0.7796	0.7000	0.7000
	Isolation Forest	0.9872	0.7215	0.1301	1.0000
	One-Class SVM	0.8670	0.5460	0.3444	0.5000
	DG-MA	0.74	0.55	0.55	0.51
Musk	LOF	1.0000	1.0000	0.9744	1.0000
	Isolation Forest	0.9553	0.4836	0.2375	1.0000
	One-Class SVM	1.0000	1.0000	0.5846	1.0000
	DG-MA	1.00	1.00	1.00	1.00

6 Conclusions

This paper investigated the Gravitation Model (GM) and its density-guided extension (DG-MA) for binary outlier detection. The proposed method integrates global centroid-based classification with local density information through a single mass adaptation step. The results show that DG-MA improves over GM in datasets with heterogeneous local structure, while preserving interpretability and linear-time inference complexity. At the same time, both models remain effective in well-separated settings, whereas purely density-based methods may be preferable for highly irregular data distributions. Overall, DG-MA provides a favourable trade-off between detection performance, transparency, and computational efficiency by translating local structural inconsistencies into global mass adjustments. These findings suggest that DG-MA is particularly suitable for data with non-uniform density structure and moderate separability, where neither purely global nor purely local methods are sufficient.

The main limitation of the proposed method is its dependence on DBSCAN parameterisation during training, as suboptimal choices of ε and $minPts$ may affect the quality of the extracted structural statistics and the resulting mass adaptation. While DG-MA reduces sensitivity through class-level aggregation, performance may degrade for highly irregular or poorly separable data distributions.

Future work will include a systematic empirical analysis of execution time, the trade-off between computational cost and detection performance, sensitivity to DBSCAN parameter choices, and soft cluster-to-class weighting strategies for heterogeneous clusters.

References

1. Han, J., Kamber, M., Pei, J.: Data Mining: Concepts and Techniques. 3 edn. Morgan Kaufmann (2011)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31

2. Peng, L., Yang, B., Chen, Y., Abraham, A.: Data gravitation based classification. Information Sciences **179**(6) (2009) 809–819
3. Shafigh, P., Hadi, S.Y., Sohrab, E.: Gravitation based classification. Information sciences **220** (2013) 319–330
4. Liu, C., Wang, W., Tu, G., Xiang, Y., Wang, S., Lv, F.: A new centroid-based classification model for text categorization. Knowledge-Based Systems **136** (2017) 15–26
5. Ester, M., Kriegel, H.P., Sander, J., Xu, X.: Density-based spatial clustering of applications with noise. In: Int. Conf. knowledge discovery and data mining. Volume 240. (1996)
6. Hawkins, D.M.: Identification of outliers. Volume 11. Springer (1980)
7. Aggarwal, C.C.: Outlier Analysis. 2nd edn. Springer (2017)
8. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM computing surveys (CSUR) **41**(3) (2009) 1–58
9. Barnett, V., Lewis, T.: Outliers in Statistical Data. 3rd edn. Wiley (1994)
10. Ramaswamy, S., Rastogi, R., Shim, K.: Efficient algorithms for mining outliers from large data sets. In: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data. (2000) 427–438
11. Breunig, M.M., Kriegel, H.P., Ng, R.T., Sander, J.: Lof: identifying density-based local outliers. In: Proceedings of the 2000 ACM SIGMOD international conference on Management of data. (2000) 93–104
12. Tang, C., He, Z., Kao, B.K.: Enhancing effectiveness of outlier detections for low density patterns. Advances in Knowledge Discovery and Data Mining (PAKDD) (2002) 535–548
13. Schölkopf, B., Platt, J.C., Shawe-Taylor, J., Smola, A.J., Williamson, R.C.: Estimating the support of a high-dimensional distribution. Neural Computation **13**(7) (2001) 1443–1471
14. Sakurada, M., Yairi, T.: Anomaly detection using autoencoders with nonlinear dimensionality reduction. In: Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis. (2014) 4–11
15. Zimek, A., Schubert, E., Kriegel, H.P.: Ensembles for unsupervised outlier detection: challenges and research questions a position paper. ACM SIGKDD Explorations Newsletter **15**(1) (2014) 11–22
16. Peng, L., Zhang, H., Yang, B., Chen, Y.: A new approach for imbalanced data classification based on data gravitation. Information Sciences **288** (2014) 347–373
17. Peng, L., Zhang, H., Zhang, H., Yang, B.: A fast feature weighting algorithm of data gravitation classification. Information Sciences **375** (2017) 54–78
18. Schubert, E., Sander, J., Ester, M., Kriegel, H.P., Xu, X.: Dbscan revisited, revisited: why and how you should (still) use dbscan. ACM Transactions on Database Systems (TODS) **42**(3) (2017) 1–21
19. Rayana, S.: Cardiotocography dataset. http://odds.cs.stonybrook.edu/cardiotocography-dataset/ (2016)
20. Rayana, S.: Japanese vowels dataset. http://odds.cs.stonybrook.edu/japanese-vowels-data/ (2016)
21. Kudo, M., Toyama, J., Shimbo, M.: Japanese vowels data set. UCI Machine Learning Repository (1999)
22. Rayana, S.: Musk dataset. http://odds.cs.stonybrook.edu/musk-dataset/ (2016)
23. Rayana, S.: Odds library. (2016)
24. Koncar, P.: Synthetic dataset for outlier detection (2018)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29924-6_31