




Revisiting Silhouette Aggregation

John Pavlopoulos^{1,2} , Georgios Vardakas³ , and Aristidis Likas³ 

¹ Department of Informatics, Athens University of Economics and Business,
Patission 76, 104 34 Athens, Greece

annis@aueb.gr

² Archimedes/Athena RC, Athens, Greece

³ Department of Computer Science and Engineering, University of Ioannina,
45110 Ioannina, Greece

g.vardakas@uoi.gr, arly@cs.uoi.gr

Abstract. Silhouette coefficient is an established internal clustering evaluation measure that produces a score per data point, assessing the quality of its clustering assignment. To assess the quality of the clustering of the whole dataset, the scores of all the points in the dataset are typically (micro) averaged into a single value. An alternative path, however, that is rarely employed, is to average first at the cluster level and then (macro) average across clusters. As we illustrate in this work with a synthetic example, the typical micro-averaging strategy is sensitive to cluster imbalance while the overlooked macro-averaging strategy is far more robust. By investigating macro-Silhouette further, we find that uniform sub-sampling, the only available strategy in existing libraries, harms the measure's robustness against imbalance. We address this issue by proposing a per-cluster sampling method. An empirical analysis on eight real-world datasets in two clustering tasks reveals the disagreement between the two coefficients for imbalanced datasets.

Keywords: cluster analysis · Silhouette · cluster validity index

1 Introduction

The silhouette coefficient [1] serves as a widely used measure for assessing the quality of clustering assignments of individual data points. It produces scores on a scale from -1 to 1 reflecting poor to excellent assignments, respectively. In real world applications, where it is widely accepted [2–4], it is common practice to average these scores to derive a single (micro-averaged) value for the entire dataset. This is the originally proposed aggregation strategy [1] and the only implementation in the popular **SCIKIT-LEARN** machine learning library in Python. An alternative aggregation strategy, however, is to average the Silhouette scores per cluster and then (macro) average across clusters, but our exploration of the related literature shows that this is a strategy that is rarely used in published studies. This is an alarming finding, because micro-averaging, e.g., in a classification context, is known to be sensitive to class imbalance [5, 6].

In this work, we focus on the effect of micro-averaging to a very well known internal cluster validity index, addressing the following research question **RQ1**: *Is micro-averaging, which is the typical strategy to aggregate Silhouette scores in cluster analysis, sensitive to cluster imbalance?* We answer this question using synthetic data, showing that micro-averaging Silhouette can produce misleading results for clustering solutions with imbalanced clusters. By contrast, we show that macro-averaging, which is rarely used in literature, is considerably more robust to this issue, because it assigns equal weight per cluster while disregarding its size.

In cluster analysis, Silhouette scores are often subsampled before being aggregated to yield a single score for a large dataset. This is a particularly useful step for computationally expensive tasks, as for example when assessing clustering solutions for a varying number of clusters to select the optimal [7]. By evaluating existing libraries in Python and R, we observe that only uniform sampling is implemented, which makes us focus on a second research question **RQ2**: *Is uniform sampling suitable when macro-averaging or is its robustness against cluster-imbalance put at stake?* The answer is the latter. Theoretically, in an extremely imbalanced dataset, the smallest cluster could even disappear when sub-sampling uniformly, which would exclude one (equal) factor from the macro-average. We address this issue, by proposing a novel per-cluster sampling strategy, which we show that it best suits macro-averaging of Silhouette scores.

Overall, the contributions of this work are:

- We compare two aggregation strategies that can be used to compute a Silhouette score for a dataset, showing that the typical micro-averaging strategy is problematic for imbalanced datasets.
- We introduce a per-cluster sampling strategy, which should be the one used along with macro-averaging.
- We quantify the sensitivity of micro-averaged Silhouette on imbalanced synthetic data, and we analyse two real-world imbalanced datasets on which the macro average should be preferred.

The remainder of this study comprises the related work (Sect. 2), a description of the Silhouette Coefficient (Sect. 3) and the aggregating strategies (Sect. 4), followed by an investigation on synthetic data (Sect. 5.1), an experimental study on real-world data (Sect. 5), and closing with remarks and future directions. Our code is publicly available in <https://github.com/ipavlopoulos/revisiting-silhouette-aggregation>.

2 Related Work

The Typical Approach. The vast majority of published studies employ micro-averaging to report the Silhouette Coefficient. In [8–10], the authors focus on the number of clusters estimation problem. In [11], they use the term Average Silhouette Width to refer to the micro-averaged Silhouette score. The authors of [12] explicitly report the implementation of SCIKIT-LEARN that employs the

micro-averaging strategy. The authors of [13], proposed a clustering algorithm that divides recursively the clustered dataset based on the maximization of the silhouette index. Similarly with the rest studies, they used the micro-average strategy, which they called as summation of silhouettes.

Exceptions to the Rule. In [14], the author observes that SPSS and R employ different implementations of the Silhouette score. They note that the latter is using micro-averaging and is the correct out of the two. We observe, however, that other libraries (packages) do not necessarily follow this paradigm. Cluster-Crit,¹ for example, compute cluster mean silhouette scores that they average to yield the final index, but this is in fact a macro-averaged score. Without any study in published literature to assess the two strategies, we argue that this is considerably problematic, because, as we show (Sect. 5), results reported with the two strategies on the same data may not be comparable. Notably, we could only detect just one study using (and explicitly stating) the macro-averaging implementation [15].

Filling the Gap. Micro-averaging is the typical and widely-used approach when aggregating Silhouette, with macro-averaging being considerably overlooked in literature of cluster analysis. Absent in existing literature is also a comparative study between the two aggregation strategies, a gap that is being bridged for classification tasks [5,6]. Our study of existing macro-averaging implementations (ClusterCrit) reveals that only uniform sampling is employed, often used for the application on large datasets. We observe, however, that uniform sampling cancels the benefits of macro-averaging (i.e., in cluster-imbalanced spaces), because the measure reduces effectively to micro-averaging. We address this gap by proposing an alternative sampling approach, well-suited to macro-averaging.

3 The Silhouette Coefficient

Data clustering is one of the most fundamental unsupervised learning tasks with numerous applications in computer science, among many other scientific fields [16,17]. Although a strict definition of clustering may be difficult to establish, a more flexible interpretation can be stated as follows: *Clustering is the process of partitioning a set of data points into groups (clusters), such that points of the same group share “common” characteristics while “differing” from points of other groups.* Data clustering can reveal the underlying data structure and hidden patterns in the data. At the same time, it is a task that poses several challenges due to the absence of labels [18], including the evaluation of clustering solutions.

Assessing the quality of a clustering solution ideally requires human expertise [19]. However, finding human evaluators could be hard, expensive and time-consuming (or even impossible for very large datasets). An alternative approach is to use clustering evaluation measures, which can be either external (supervised) or internal (unsupervised) [20]. The former, as the name suggests,

¹ <https://cran.r-project.org/web/packages/clusterCrit/vignettes/clusterCrit.pdf>.

use external information (e.g., classification labels) as the ground truth cluster labels. Well known external evaluation measures are Normalised Mutual Information (NMI) [21], Adjusted Mutual Information (AMI) [22], Adjusted Rand Index (ARI) [23,24], etc. External information, however, is not typically available in real-world scenarios. In such cases we resort to internal evaluation measures, which are solely based on information intrinsic to the data. Although other internal evaluation measures have been proposed [25,26], we focus on the most commonly-employed, and successful one [27], which is the silhouette coefficient [1].




Fig. 1. Illustration of the elements involved in the computation of the silhouette score $s(x_i)$ for a given data point x_i that belongs to cluster C_I .

The silhouette coefficient [1] is a measure to assess clustering quality, which does not depend on external knowledge and that does not require ground truth labels. A good clustering solution, according to this measure, assumes compact and well-separated clusters. Formally, given a dataset $X = \{x_1, \dots, x_N\}$ that is partitioned by a clustering solution $f : X \rightarrow \{C_1, \dots, C_K\}$ into K clusters, the silhouette coefficient for point $x_i \in X$ is based on two values, the inner and the outer cluster distance. The former, denoted as $a(x_i)$, is the average distance between x_i and all other points within the cluster C_I that x_i belongs to (i.e., $f(x_i) = C_I$):

$$a(x_i) = \frac{1}{|C_I| - 1} \sum_{x_j \in C_I, i \neq j} d(x_i, x_j), \quad (1)$$

where $|C_I|$ represents the cardinality of cluster C_I and $d(x_i, x_j)$ is the distance between x_i and x_j . The $a(x_i)$ value quantifies how well the point x_i fits within its cluster. For example, in Fig. 1, $a(x_i)$ measures the average distance of x_i to the points in its cluster C_I . A low value of $a(x_i)$ indicates that x_i is close to the other members of that cluster, suggesting that x_i is probably grouped correctly. Conversely, a higher value of $a(x_i)$ indicates that x_i is not well-placed in that cluster. In addition, the silhouette score requires the calculation of the minimum

average outer-cluster distance $b(x_i)$ per point x_i , defined as:

$$b(x_i) = \min_{C_J \neq f(x_i)} \frac{1}{|C_J|} \sum_{x_j \in C_J} d(x_i, x_j). \quad (2)$$

A large $b(x_i)$ value indicates that x_i significantly differs from the points of the closest cluster. In Fig. 1, the closest cluster (which minimises $b(x_i)$) is C_J . Considering both $a(x_i)$ and $b(x_i)$, the silhouette score of x_i is defined as:

$$s(x_i) = \frac{b(x_i) - a(x_i)}{\max\{a(x_i), b(x_i)\}}. \quad (3)$$

It is evident that the silhouette score $s(x_i)$, defined in Eq. 3 for a data point x_i , falls within the range $-1 \leq s(x_i) \leq 1$. Values close to 1 indicate that point x_i belongs to a compact, well-separated group. In contrast, values close to -1 suggest that another cluster assignment for that point would have been a better option.

4 Methods

The Silhouette Coefficient provides a score that grades the cluster assignment of a data point. To obtain a single score for all the points $x \in X$, the typical approach is micro-averaging (see Sect. 2) that averages all the individual scores. The alternative macro-averaging approach averages first the scores per cluster, and then (macro) average the latter.

4.1 Micro-averaged Silhouette: *The Typical Index*

Micro-averaging silhouette at the point level (sample mean) is defined as follows:

$$S_{micro}(X) = \frac{1}{N} \sum_{x_i \in X} s(x_i). \quad (4)$$

This is the originally proposed averaging strategy by the study that introduced silhouette [1], the one adopted by **SCIKIT-LEARN**, and it is the typical approach employed in literature [11]. However, we show in Sect. 5 that it is not effective in the case of imbalanced clusters, which is very common in real-world datasets.

4.2 Macro-averaged Silhouette: *The Overlooked Index*

When clusters are perfectly balanced, the sample mean is a reasonable aggregation strategy. The assumption of perfectly balanced clusters, however, cannot be guaranteed in the real world, where clusters are often imbalanced. In such cases, and when small clusters also matter (e.g., diagnostic reports about rare medical diseases), we argue that micro-averaging is not effective while macro-averaging is robust. This issue is known in fields such as supervised learning [5], but it has

not been studied yet for clustering. To compute the macro-Silhouette, a score S_c is computed for each cluster C_i as follows:

$$S_C(C_i) = \frac{1}{|C_i|} \sum_{x_i \in C_i} s(x_i). \quad (5)$$

This score measures how compact and well separated a cluster is given a clustering solution. For K clusters in that solution, we end up with a set of K cluster silhouette values $\{S_{C_1}, \dots, S_{C_K}\}$. The average of these K scores, defined as the macro-averaged Silhouette, can be used to assess the dataset clustering and is more formally defined as follows:

$$S_{macro}(X) = \frac{1}{K} \sum_{k=1}^K S_C(C_i). \quad (6)$$

We note that macro-averaging assumes equal weight between the clusters, but other approaches also apply. The weighted average, for example, where weights reflect the support (i.e., the number of points per cluster, normalised), is closer to micro-averaging in nature. Furthermore, other statistics could be applied, such as the max (or the min), capturing the most (least) compact and well (bad) separated cluster.

4.3 Per-cluster Sampling: *Efficient and Robust Macro-Silhouette*

The computation of the silhouette coefficient for all the N points in a dataset requires the computation of a pairwise distance matrix at the cost of $\mathcal{O}(N^2)$ operations. This is demanding in terms of computational and space complexity and, hence, not scalable for large datasets [28]. The typical approach to tackle this problem is to compute the silhouette score using a uniformly selected subsample of the dataset.

In a cluster-imbalanced problem, the typical (uniform across data points) sampling may favour the major cluster and may even disregard completely one of the minor clusters. We argue that this practice contradicts the nature of macro-averaging, which assumes that clusters are equally weighted when averaged. To solve this problem, when macro-averaging is aimed, we propose that sampling takes place per cluster, following the macro-averaging spirit.

More specific, we create a subsample of size L for computing the macro-averaged silhouette score by uniformly selecting a subset of L/K points from each cluster C_i , where $i = 1, \dots, K$. In this way, we ensure that all the clusters contribute a sufficient number of data points to the subsample, preserving the robustness of macro-Silhouette.

5 Experiments

5.1 Analysis on Synthetic Data

We created a synthetic dataset consisting of four Gaussian clusters, each with 100 points and a variance of 0.1, shown on the left of Fig. 2. The micro- and




Fig. 2. Synthetic dataset, shown on the left, with four equibalanced clusters. The same space is shown on the right, but the relatively distant cluster B now comprises 5,000 points more, yielding a heavy cluster-imbalance. Silhouette is reported per dataset per aggregation strategy. Micro-averaging increases unreasonably by a large margin.

macro-averaged Silhouette scores are both 0.623, a score that is far from perfect due to the overlapping clusters C and D. The same space is shown on the right, but the lower distant cluster B is now populated with 5,000 points more, resulting in a heavy cluster imbalance. Specifically, the ratio of the smallest to the largest cluster for this dataset is $r = 0.02$.

When moving from the (balanced) space on the left to the (imbalanced) space on the right, micro changes considerably. That is because the points added to the distant cluster B directly influence the point-level average value upwards (i.e., a relative increase of 41%). The overlap of the minor clusters C and D is less important in this case, when compared to the well-separated major cluster B. By contrast, macro-average remains robust to this change, because each cluster is equally weighted in the average, disregarding their size. As is shown in Fig. 3, the sensitivity of micro-averaging to imbalance becomes apparent early on and it could continue increasing if we continued adding points.




Fig. 3. Silhouette score of the dataset of Fig. 2, micro- and macro-averaged, for a varying number of points added.

Per-cluster Sampling. Using the same dataset, we assess the robustness of the proposed per-cluster sampling. This is clearly shown in Fig. 4, by comparing the

typical uniform and the proposed per-cluster sampling of the macro-averaged Silhouette score. The more the imbalance, as we move to the right of the Figure, the more the fluctuations of the macro-averaged Silhouette score computed on a uniform sample of 100 points. The per-cluster sampling, on the other hand, remains robust. Similar fluctuations are observed on uniform sampling and micro-averaging (in red).




Fig. 4. Macro-averaged Silhouette score, computed on uniform and per-cluster samples of 100 points, as the size of cluster B of Fig. 2 increases. (Color figure online)

Major Overlapping Cluster. When adding points to a distant cluster, we observe an increase of micro-averaged Silhouette as opposed to its macro-averaged counterpart. The situation is different, however, when the cluster we add points to is close to others. Figure 5 depicts this space, where we observe that micro-averaged Silhouette drops in value in this case. Macro-Silhouette, on the other hand, again remains robust. This is mainly because we add points that overlap with nearby clusters, reducing the overall score instead of increasing it, as was the case when we added points to a distant cluster (Fig. 3).




Fig. 5. Synthetic dataset, equilibrated as in Fig. 2 on the left, but now cluster B (to which we add 5,000 points on the right) is very close to clusters D and A.

Estimating k (the Silhouette Method). Silhouette has been suggested as an alternative to the problematic “elbow” method when estimating the number of

clusters [7]. That is, the quality of a clustering solution (e.g., the predictions of KMeans) is evaluated for different numbers of clusters. The number that leads to the highest Silhouette score is chosen as the optimal one. We applied this method on a synthetic imbalanced dataset of four isotropic Gaussian blobs,² three with 100 points and one with 2,300. By undersampling from the major cluster, we also produce an equibalanced version of this space where clusters comprised 100 points each. By applying KMeans, then, with k ranging from 2 to 10, we measured the micro and macro averaged Silhouette per space per k , using uniform (for micro) and per-cluster (for macro) sampling of 100 points. As is shown in Fig. 6, macro-averaged Silhouette reaches a maximum (blue star) on the ground truth number of clusters (red line) in the imbalanced dataset. Micro-averaged Silhouette is maximised for a different number of clusters. In the undersampled (balanced) version of this dataset, shown on the right, both strategies reach their maximum on the ground truth number of clusters, i.e., four, where the red vertical line is. This result can be explained by the lack of robustness of uniform-sampled micro-averaged Silhouette (see Fig. 4), which is deceiving when estimating the number of clusters. The robust per-cluster-sampled macro-averaged Silhouette, on the other hand, is not affected.

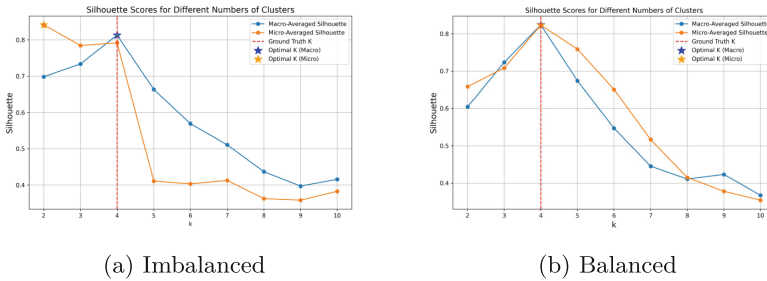


Fig. 6. Estimating the optimal number of clusters (shown with star) when using micro (orange) and macro (blue) averaged Silhouette. We use both, an imbalanced dataset of four isotropic Gaussians, and an undersampled (balanced) version of the same space. A red vertical line shows the ground truth number of clusters. (Color figure online)

5.2 Application on Real-World Datasets

We employed eight real-world datasets [29] of various types (numeric, time-series, images), sizes (from 150 to more than 500,000 items), dimensionality, and with a varying cluster imbalance. To estimate the latter, we computed the ratio (r) of the size of the smallest to the largest cluster. Table 1 displays these datasets,

² https://scikit-learn.org/stable/modules/generated/sklearn.datasets.make_blobs.html.

sorted according to their imbalance, which ranged from high ($r = 0.03$) to low ($r = 1.00$).

These eight datasets are summarised below:

- PENDIGITS comprises 10,992 pen-based handwritten digits (from 0 to 9). Each data item is represented by a 16-dimensional vector containing pixel coordinates. DIGITS also comprises (1,797) images of handwritten digits, but each item is an image of 8×8 pixels, resulting in $d = 64$ features.
- COVER TYPE contains cartographic variables for predicting forest cover types. The dataset includes 581,012 samples and 54 features, such as elevation, aspect, slope, and soil type. The cover types are classified into seven categories and it is a highly imbalanced dataset ($r = 0.03$).
- GAS SENSOR consists of measurements from 16 chemical sensors exposed to different gases over a period of several months. The dataset includes 13,910 samples and 128 features. It is used for studying the drift in sensor responses over time and developing algorithms for sensor calibration.
- WINE contains the results of a chemical analysis of wines grown in the same region in Italy but derived from three different cultivars. The analysis determined the quantities of 13 constituents found in each of the three types of wines.
- IRIS comprises the lengths and widths of the sepals and petals of iris flowers.
- GLASS contains the chemical compositions of glass samples, which are classified into seven types of glass.
- TCGA is a collection of gene expression profiles obtained from RNA sequencing of various cancer samples. It includes 801 data instances, clinical information, normalised counts, gene annotations, and 6 cancer types' pathways.
- MICE consists of the expression levels of 77 proteins/protein modifications that produced detectable signals in the nuclear fraction of the cortex. It includes 1,080 data points and 8 eight classes of mice based on the genotype, behaviour, and treatment characteristics.

Evaluation Metrics. We compute micro and macro Silhouette scores, as internal validation measures. We also employ external validation measures that use the ground truth labels to assess the clustering solution, presented for completeness. The normalized mutual information (NMI) score measures the similarity between two clusterings by normalizing the mutual information score [30]. A higher score indicates a better match between the cluster labels and the ground truth labels. The adjusted mutual information (AMI) adjusts MI for chance groupings [30]. It measures the agreement between two clustering assignments and is normalized against the entropy of the labels to yield a score between 0 and 1. The adjusted rand index (ARI) measures the similarity between two data clusterings by adjusting Rand Index (RI) to account for the chance grouping of elements [23]. The score ranges from -1 to 1 , with higher values indicating better performance.

Table 1. Real-world datasets of varying dimensionality (d), size (N), number of clusters (k), imbalance ratio of smallest to largest cluster (r). The average macro (MaS) and micro (MiS) Silhouette score, sampled uniformly and per-cluster respectively, is reported across three runs (st. error of mean), along with NMI, ARI, AMI. Sorted by r .

Dataset	Type	N	d	k	r	MaS	MiS	NMI	ARI	AMI
Iris	Numeric	150	4	3	1.00	0.46 ± 0.01	0.46 ± 0.01	0.66	0.62	0.66
Digits	Image	1797	64	10	0.95	0.11 ± 0.01	0.13 ± 0.01	0.69	0.56	0.69
Pendigits	Time-series	10992	16	10	0.92	0.24 ± 0.01	0.23 ± 0.01	0.67	0.53	0.67
Mice Protein	Numeric	1080	77	8	0.70	0.13 ± 0.01	0.12 ± 0.01	0.26	0.14	0.25
Wine	Numeric	178	13	3	0.68	0.30 ± 0.01	0.29 ± 0.01	0.86	0.88	0.86
Gas Sensor	Time-series	13910	128	6	0.55	0.22 ± 0.03	0.27 ± 0.01	0.19	0.07	0.19
Glass	Numeric	214	9	6	0.12	0.20 ± 0.01	0.31 ± 0.01	0.32	0.17	0.29
Covertypes	Numeric	110393	54	7	0.03	0.25 ± 0.01	0.08 ± 0.01	0.13	0.05	0.13

Experimental Settings. Missing values in datasets were replaced with the mean value of the respective feature.³ We standardized all the features per dataset, by removing the mean and scaling to unit variance,⁴ to avoid numerical instabilities in the computations [31]. We trained KMeans per dataset, using the ground truth number of clusters and selecting initial cluster centroids with KMeans++.

Results with Fixed Ground Truth Number of Clusters. Table 1 presents the evaluation results of KMeans across datasets, setting k to the ground truth number of clusters of each dataset. The external validation scores (NMI, ARI, AMI) are high for balanced datasets ($r > 0.9$), they drop in two out of three mild-imbalanced ones ($0.5 \leq r \leq 0.7$), and they are overall low for both highly imbalanced datasets. When computing Silhouette, we employ uniform (for micro) and per-cluster (for macro) sampling of 100 points per dataset, repeating three times and reporting the average and the standard error of the mean. We observe that the two indices are very close to each other in the balanced and mild imbalanced zones, with any differences not exceeding the standard error of the mean. By contrast, in the highly imbalanced zone, the two indices are far from each other. The sensitivity of micro-averaging to cluster imbalance (Sect. 5.1) can explain the observed difference between the two indices. When MiS is greater (e.g., GLASS), a major distant cluster may be present while when MaS is greater than MiS (e.g., COVER TYPE), a major cluster may be close to other ones.

³ <https://scikit-learn.org/stable/modules/generated/sklearn.impute.SimpleImputer.html>.

⁴ <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.StandardScaler.html>.

Results on k -Estimation. We experimented also with estimating k using the Silhouette method. That is by applying clustering for various k values and assessing the two Silhouette aggregation strategies regarding their ability to yield a maximum score for the ground truth number of clusters. The ground truth number of clusters is shown with k in Table 1 and depicted with a red vertical line in Fig. 7. In Fig. 7, we see that micro and macro averaging yield the same optimal k in five out of eight datasets, all balanced (IRIS, PENDIGITS) or mildly imbalanced (MICE, WINE, GAS SENSOR). In the heavily imbalanced datasets of GLASS and COVER TYPE, as expected, the optimal k differs. Overall, the macro-average yields an optimal k that is the same as the ground truth k in two datasets (DIGITS and WINE) while typical micro-averaging yields an optimal k that is the same as the grand truth in one (WINE).

6 Discussion

Silhouette Aggregation and Dataset. As was shown in Sect. 5, the typical micro-averaged Silhouette score is vulnerable to cluster imbalance while the rarely-used macro-averaged Silhouette is far more robust. This means that when there is indication of an imbalanced dataset, it is macro-averaged Silhouette that should be trusted for evaluating clustering solutions. This is the case for the two most imbalanced real-world datasets used in this work. GLASS achieves a lower macro-average compared to micro, which could be explained by a major distant cluster, as in Fig. 2. COVER TYPE, on the other hand, achieves a higher macro-average, compared to micro, which fits the synthetic example of Fig. 5 with a major cluster being nearby others.

Silhouette Aggregation Per Domain. The choice of the aggregation strategy depends also on the application domain. In predictive maintenance, for example, major faults are of much higher importance compared to rare events, because an accurate estimation may allow better logistics and administration for the company (e.g., early orders, select appropriate workstations, etc.). In this case, micro-Silhouette should be the selected index, if clustering was applied. On the other hand, biomedical clustering applications would likely select macro-Silhouette, because rare medical conditions exist (e.g., adverse drug events, etc.) and should not be considered of less importance to frequently occurring ones.

Appropriate Sampling. During our study of related work and existing implementations (Sect. 2), we observed that the only sampling strategy implemented was uniform. This strategy, however, is not appropriate for the macro-averaged Silhouette (Sect. 4.3). Therefore, we proposed a per-cluster sampling strategy, which we showed that it is considerably more robust compared to standard uniform sampling. This contribution can be particularly important for big datasets, because computing Silhouette is $\mathcal{O}(N^2)$. As was shown in Fig. 4, per-cluster sampling is robust to imbalance and yields approximately the same score even when the subsampled space is 2% of the original (i.e., rightmost of Fig. 4).

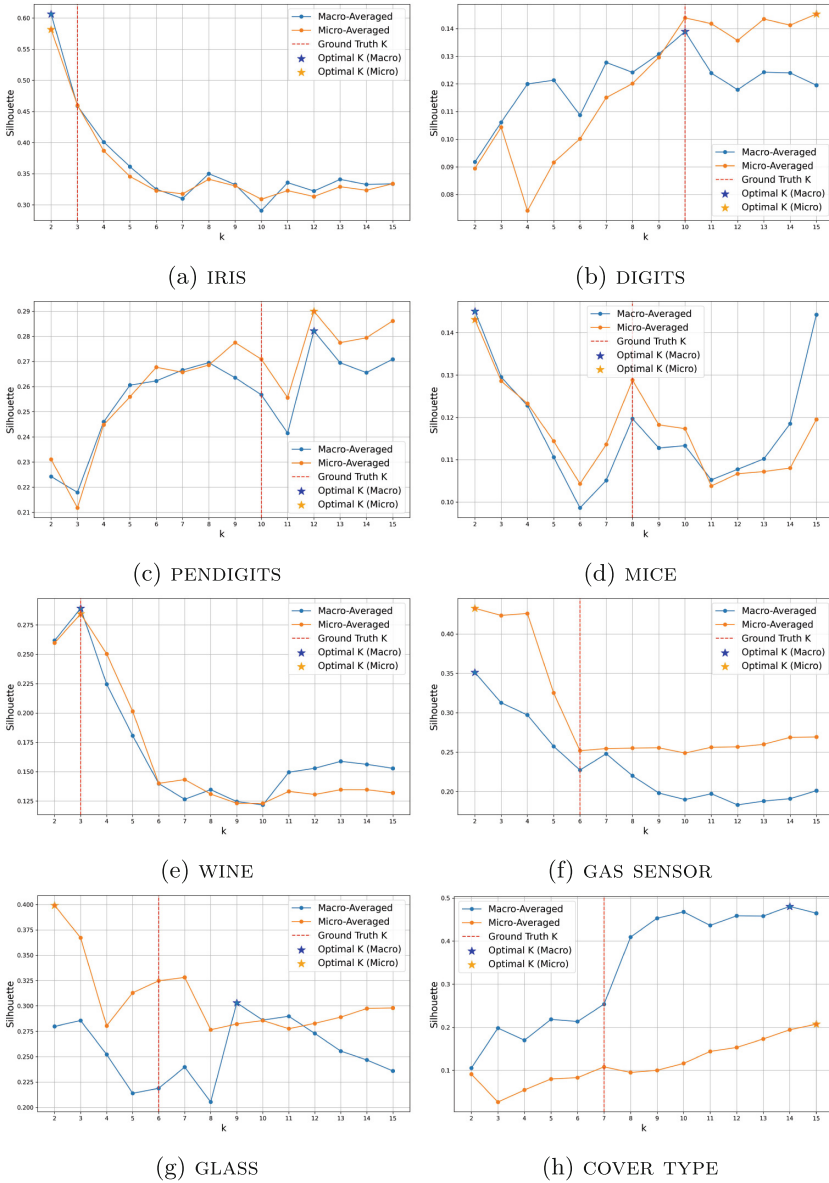


Fig. 7. Micro (blue) and macro (blue) Silhouette per dataset. Clustering produced with KMeans for varying k to select the optimal number of clusters (shown with a star). The ground truth number of clusters is shown with a vertical red line. (Color figure online)

7 Conclusions

This study shows that, although heavily overlooked, macro-Silhouette is a serious option that should be considered in cluster analysis. By focusing on subsampling, an important step for large datasets, we find that standard uniform sampling is not appropriate for macro-Silhouette. Hence, we propose a novel robust per-cluster sampling strategy that follows in nature the macro-Silhouette computation. By employing eight real-world datasets of varying cluster-imbalance, we undertake an analysis showing that the two indices disagree in imbalanced datasets.

Acknowledgements. – This work has been supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

– This research project is implemented in the framework of H.F.R.I. call “Basic research Financing (Horizontal support of all Sciences)” under the National Recovery and Resilience Plan “Greece 2.0” funded by the European Union - NextGenerationEU (H.F.R.I. ProjectNumber: 15940).

References

1. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20**, 53–65 (1987)
2. Layton, R., Watters, P., Dazeley, R.: Evaluating authorship distance methods using the positive silhouette coefficient. *Nat. Lang. Eng.* **19**(4), 517–535 (2013)
3. Bafna, P., Pramod, D., Vaidya, A.: Document clustering: TF-IDF approach. In: 2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT), pp. 61–66. IEEE (2016)
4. Tambunan, H.B., Barus, D.H., Hartono, J., Alam, A.S., Nugraha, D.A., Usman, H.H.H.: Electrical peak load clustering analysis using k-means algorithm and silhouette coefficient. In: 2020 International Conference on Technology and Policy in Energy and Electric Power (ICT-PEP), pp. 258–262. IEEE (2020)
5. Gaudreault, J.-G., Branco, P.: Empirical analysis of performance assessment for imbalanced classification. *Mach. Learn.* 1–43 (2024)
6. Suhaimi, N.S., Othman, Z., Yaakub, M.R.: Comparative analysis between macro and micro-accuracy in imbalance dataset for movie review classification. In: Proceedings of Seventh International Congress on Information and Communication Technology: ICICT 2022, London, Volume 3, pp. 83–93. Springer, Cham (2022)
7. Schubert, E.: Stop using the elbow criterion for k-means and how to choose the number of clusters instead. *ACM SIGKDD Explorations Newsl.* **25**(1), 36–42 (2023)
8. Azimi, R., Ghayekhloo, M., Ghofrani, M., Sajedi, H.: A novel clustering algorithm based on data transformation approaches. *Expert Syst. Appl.* **76**, 59–70 (2017)
9. Dudek, A.: Silhouette index as clustering evaluation tool. In: Jajuga, K., Batóg, J., Walesiak, M. (eds.) SKAD 2019. SCDAKO, pp. 19–33. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-52348-0_2
10. Ünlü, R., Xanthopoulos, P.: Estimating the number of clusters in a dataset via consensus clustering. *Expert Syst. Appl.* **125**, 33–39 (2019)

11. Batool, F., Hennig, C.: Clustering with the average silhouette width. *Comput. Stat. Data Anal.* **158**, 107190 (2021)
12. Shahapure, K.R., Nicholas, C.: Cluster quality analysis using silhouette score. In: 2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA), pp. 747–748. IEEE (2020)
13. Kang, J.H., Park, C.H., Kim, S.B.: Recursive partitioning clustering tree algorithm. *Pattern Anal. Appl.* **19**, 355–367 (2016)
14. Řezanková, H.: Different approaches to the silhouette coefficient calculation in cluster evaluation. In: 21st International Scientific Conference AMSE Applications of Mathematics and Statistics in Economics, pp. 1–10 (2018)
15. Brun, M., et al.: Model-based evaluation of clustering validation measures. *Pattern Recognit.* **40**(3), 807–824 (2007)
16. Jain, A.K.: Data clustering: 50 years beyond k-means. *Pattern Recognit. Lett.* **31**(8), 651–666 (2010)
17. Ezugwu, A.E., et al.: A comprehensive survey of clustering algorithms: state-of-the-art machine learning applications, taxonomy, challenges, and future research prospects. *Eng. Appl. Artif. Intell.* **110**, 104743 (2022)
18. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: a review. *ACM Comput. Surv. (CSUR)* **31**(3), 264–323 (1999)
19. Von Luxburg, U., Williamson, R.C., Guyon, I.: Clustering: science or art? In: Proceedings of ICML Workshop on Unsupervised and Transfer Learning, pp. 65–79. JMLR Workshop and Conference Proceedings (2012)
20. Rendón, E., Abundez, I., Arizmendi, A., Quiroz, E.M.: Internal versus external cluster validation indexes. *Int. J. Comput. Commun.* **5**(1), 27–34 (2011)
21. Estévez, P.A., Tesmer, M., Perez, C.A., Zurada, J.M.: Normalized mutual information feature selection. *IEEE Trans. Neural Netw.* **20**(2), 189–201 (2009)
22. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* **11**(95), 2837–2854 (2010)
23. Hubert, L., Arabie, P.: Comparing partitions. *J. Classif.* **2**(1), 193–218 (1985)
24. Chacón, J.E., Rastrojo, A.I.: Minimum adjusted rand index for two clusterings of a given size. *Adv. Data Anal. Classif.* 1–9 (2022)
25. Caliński, T., Harabasz, J.: A dendrite method for cluster analysis. *Commun. Stat.-Theory Methods* **3**(1), 1–27 (1974)
26. Davies, D.L., Bouldin, D.W.: A cluster separation measure. *IEEE Trans. Pattern Anal. Mach. Intell.* (2), 224–227 (1979)
27. Arbelaiz, O., Gurrutxaga, I., Muguerza, J., Pérez, J.M., Perona, I.: An extensive comparative study of cluster validity indices. *Pattern Recognit.* **46**(1), 243–256 (2013)
28. Capó, M., Pérez, A., Lozano, J.A.: Fast computation of cluster validity measures for bregman divergences and benefits. *Pattern Recognit. Lett.* **170**, 100–105 (2023)
29. Dua, D., Graff, C.: UCI machine learning repository (2017)
30. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: is a correction for chance necessary? In: Proceedings of the 26th Annual International Conference on Machine Learning, pp. 1073–1080 (2009)
31. Celebi, M.E., Kingravi, H.A., Vela, P.A.: A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Syst. Appl.* **40**(1), 200–210 (2013)