



HAL
open science

Selecting the Number of Clusters K with a Stability Trade-off: An Internal Validation Criterion

Alex Mourer, Florent Forest, Mustapha Lebbah, Hanane Azzag, Jérôme Lacaille

► **To cite this version:**

Alex Mourer, Florent Forest, Mustapha Lebbah, Hanane Azzag, Jérôme Lacaille. Selecting the Number of Clusters K with a Stability Trade-off: An Internal Validation Criterion. The Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), May 2023, Osaka (JP), Japan. pp.210-222, 10.1007/978-3-031-33374-3_17. hal-04456322

HAL Id: hal-04456322

<https://hal.science/hal-04456322v1>

Submitted on 30 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Selecting the Number of Clusters K with a Stability Trade-off: an Internal Validation Criterion

Alex Mourer^{*1,4}, Florent Forest^{*2,4}, Mustapha Lebbah³, Hanane Azzag², and Jérôme Lacaille⁴

¹ SAMM, Université Paris 1 Panthéon Sorbonne, France

² LIPN (CNRS UMR 7030), Université Sorbonne Paris Nord, France
f@florentfo.rest

³ DAVID lab, Université de Versailles/Paris-Saclay, France

⁴ Safran Aircraft Engines, France

Abstract. Model selection is a major challenge in non-parametric clustering. There is no universally admitted way to evaluate clustering results for the obvious reason that no ground truth is available. The difficulty to find a universal evaluation criterion is a consequence of the ill-defined objective of clustering. In this perspective, clustering stability has emerged as a natural and model-agnostic principle: an algorithm should find stable structures in the data. If data sets are repeatedly sampled from the same underlying distribution, an algorithm should find similar partitions. However, stability alone is not well-suited to determine the number of clusters. For instance, it is unable to detect if the number of clusters is too small. We propose a new principle: a good clustering should be stable, and within each cluster, there should exist no stable partition. This principle leads to a novel clustering validation criterion based on between-cluster and within-cluster stability, overcoming limitations of previous stability-based methods. We empirically demonstrate the effectiveness of our criterion to select the number of clusters and compare it with existing methods. Code is available at <https://github.com/FlorentF9/skstab>.

Keywords: clustering · model selection · stability · internal validation.

1 Introduction

Clustering is an unsupervised learning technique aiming at discovering structure in unlabeled data. It can be defined as the “partitioning of data into groups (a.k.a. clusters) so that similar [...] elements share the same cluster and the members of each cluster are all similar” [6]. These goals are contradictory because of the non-transitivity of similarity: if A is similar to B , and B is similar to C , A is not necessarily similar to C . Since clustering is an ill-posed problem, it cannot be properly solved using this definition, and algorithms often optimize only one

* Equal contribution. Supported by ANRT CIFRE grants and Safran Aircraft Engines.

of its aspects. For instance, K -means only guarantees that dissimilar objects are separated, and on the other hand, single linkage clustering only guarantees that similar objects will end up in the same cluster. As a consequence, model selection is a major challenge in non-parametric clustering.

In the sample-based framework adopted in this work, model selection assesses whether partitions found by an algorithm correspond to meaningful structures of the underlying distribution, and not just artifacts of the algorithm or sampling process [44,43]. Practitioners need to evaluate clustering results in order to select the best parameters for an algorithm (e.g. the number of clusters K) or choose between different algorithms. Plenty of evaluation methods exist in literature, but they usually incorporate strong assumptions on the geometry of clusters or on the underlying distribution.

There is a need for a general, model-agnostic evaluation method. Clustering stability has emerged as a principle stating that "to be meaningful, a clustering must be both *good* and the only *good* clustering of the data, up to small perturbations. Such a clustering is called stable. Data that contains a stable clustering is said to be clusterable" [35]. Hence, a clustering algorithm should discover stable structures in the data. In statistical learning terms, if data sets are repeatedly sampled from the same underlying distribution, an algorithm should find similar partitions. As we do not have access to the data-generating distribution, perturbed data sets are obtained either by sampling or injecting noise into the original data. Stability seems to be an elegant principle, but there are still severe limitations in practice. For instance, stability does not necessarily depend on clustering outcomes but can be solely related to properties of the data such as symmetries [9]. As outlined in [51], there exist various protocols to estimate stability. Unfortunately, a thorough study that evaluates them in practice is lacking.

Contributions We propose a method for quantitatively and visually assessing the presence of structure in clustered data. The main contributions of our work can be stated as follows:

- To our knowledge, this is the first large-scale empirical study on clustering stability analysis.
- A novel definition of clustering is proposed, based on between-cluster and within-cluster stability. Based on this definition, we introduce Stadion, the stability difference criterion, along with an interpretable visualization tool, called *stability paths*.
- We show that additive noise perturbation is reliable, and a methodology to determine the amount of perturbation is proposed.
- We assess the ability of Stadion to select the number of clusters K on a vast collection of data sets and compare it with state-of-the-art methods.

2 Related work

Internal clustering indices measure the quality of a clustering when ground-truth labels are unavailable. Most criteria rely on a combination of between-cluster and

within-cluster distances. Between-cluster distance measures how distinct clusters are dissimilar, while within-cluster distance measures how elements belonging to the same cluster are similar. Unfortunately, this incorporates a prior on the geometry of clusters [18,14,15,40,39,46,17].

Stability analysis for clustering validation is a long-established technique. It can be traced back as far as 1973 [45] and from there has drawn increasing attention [11,32,9,7,8,51]. Some works concluded that stability is not a well-suited tool for model selection [43]. In the general case, stability can only detect if the number of clusters is too large for the K -means algorithm (see Figure 1). A partition with too few clusters is indeed stable, except for perfectly symmetric distributions. More accurately, these works proved that the asymptotic stability of risk-minimizing clustering algorithms, as sample size grows to infinity, only depends on whether the objective function has one or several global minima.

Albeit significant theoretical efforts, few empirical studies have been conducted. Each study focuses on specific practical implementations of stability, but as mentioned in [51,10], a thorough study comparing all protocols in practice does not exist and a more objective evaluation of these results is warranted.

3 Clustering stability

A data set $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ consists in N independent and identically distributed (i.i.d.) samples, drawn from a data-generating distribution \mathcal{P} on an underlying space \mathcal{X} . Formally, a clustering algorithm \mathcal{A} takes as input the data set \mathbf{X} , some parameter $K \geq 1$, and outputs a clustering $\mathcal{C}_K = \{C_1, \dots, C_K\}$ of \mathbf{X} into K disjoint sets. Thus, a clustering can be represented by a function $\mathbf{X} \rightarrow \{1, \dots, K\}$ assigning a label to every point of the input data set. Some algorithms can be extended to construct a partition of the entire underlying space. This partition is represented by an extension operator function $\mathcal{X} \rightarrow \{1, \dots, K\}$ (e.g. for center-based algorithms, we compute the distance to the nearest center).




Fig. 1. Example data set with three clusters. The labels correspond to the K -means clustering result for $K = 2, 3$ and 4 . K -means is stable even if K is too small.

Let \mathbf{X} and \mathbf{X}' be two data sets drawn from the same distribution and note \mathcal{C}_K and \mathcal{C}'_K their respective clusterings. Let s be a similarity measure such that

$s(\mathcal{C}_K, \mathcal{C}'_K)$ measures the agreement between the two clusterings. Then, for a given sample size N , the stability of a clustering algorithm \mathcal{A} is defined as the expected similarity between \mathcal{C}_K and \mathcal{C}'_K on different data sets \mathbf{X} and \mathbf{X}' , sampled from the same distribution \mathcal{P} ,

$$\text{Stab}(\mathcal{A}, K) := \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{P}^N} [s(\mathcal{C}_K, \mathcal{C}'_K)]. \quad (1)$$

This quantity is unavailable in practice, as we have a finite number of samples, so it needs to be estimated empirically. Various methods have been devised to estimate stability using perturbed versions of \mathbf{X} .

The first methods used in literature are based on resampling the original data set (splitting in half [45], subsampling [11], bootstrapping [19,20], jackknife [53], etc.). Another method consists in adding random noise either to the data points [36] or to their pairwise distances [49,3]. For high-dimensional data, alternatives are random projections or randomly adding or deleting variables [45]. Once the perturbed data sets are generated, there are several ways to compare the resulting clusterings. With noise-based methods, it is possible to compare the clustering of the original data set (reference clustering) with the clusterings obtained on perturbed data sets, or to compare only clusterings obtained on the latter. With sampling-based methods, we can compare overlapping subsamples on data points where both clusterings are defined [19], or compare clusterings of disjoint subsamples (using for instance an extension operator or a supervised classifier to transfer labels from one sample to another [32]). Finally, possible similarity measures include external indices such as the ARI [19,55].

Before discussing in details the mechanisms of stability, we introduce a trivial example to illustrate its main issue: it cannot detect in general whenever K is too small. Consider the example presented in Figure 1 with three clusters. On any sample from such a distribution, as soon as we have a reasonable amount of data, K -means with $K = 2$ always constructs the solution separating the left cluster from the two right clusters. Consequently, it is stable despite $K = 2$ being the wrong number of clusters. This situation was pointed out in [9].

In the case of algorithms that minimize an objective function (e.g. center-based or spectral), two different sources of instability have been identified [51]. First, *jittering* is caused by assignment changes at cluster boundaries after perturbation. Therefore, strong jitter is produced when a cluster boundary cuts through high-density regions. Second, *jumping* refers to the algorithm ending up in different local minima. The most important cause of jumping is initialization. Another cause is the presence of several global minima of the objective function. This happens if there are perfect symmetries in the distribution, which is very unlikely in real-world data. Examples are provided as supplementary material.

However, practitioners mainly use algorithms with consistent initialization strategies. For instance with K -means, we keep the best trial over a large number of runs and use the K -means++ heuristic. This initialization tends to make K -means deterministic, differently from the random initialization proposed in [51,13], which allows jumping to occur whenever $K > K^*$, where K^* is the true number of clusters. Throughout this work, we consider a setting with large

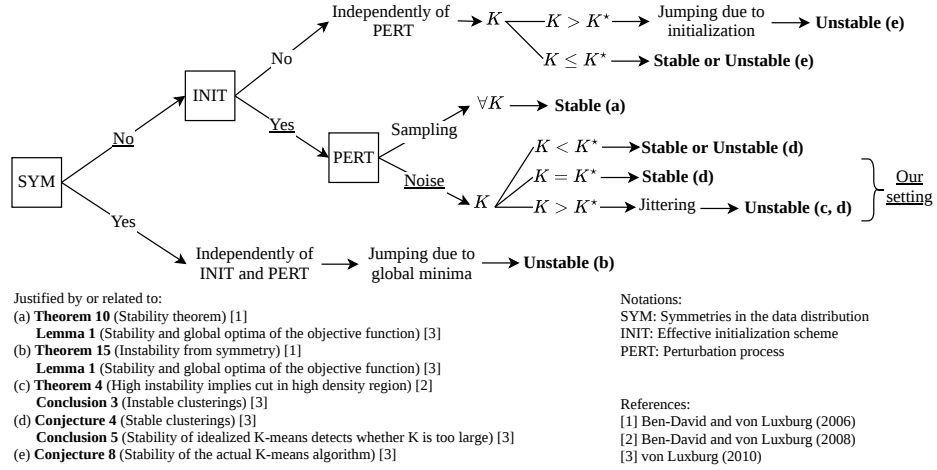


Fig. 2. Diagram explaining sources of instability in different settings, based on theoretical results for K -means, with large sample size, assuming $K \ll N$ and the underlying distribution has K^* well-separated clusters that can be represented by K -means. We consider no symmetries, effective initialization and noise-based perturbation, thus instability (due to jittering) arises when K is too large, and when K is too small whenever cluster boundaries are in high-density regions.

enough sample size, without perfect symmetries and with consistent initialization, that we deem to be realistic. Thus, we do not rely on jumping as the main source of instability even when $K > K^*$, and rather rely on jittering. As a consequence, we need a perturbation process that produces jittering. We settle for noise-based perturbation, because as soon as N is reasonably large, resampling methods become trivially stable whenever there is a single global minimum [9,51]. We summarize important results in the diagram Figure 2 and provide a simple example where sampling methods such as [11,32] fail in the supplementary material. To conclude, in our setting, a noise-based perturbation process causes jittering, enabling stability to indicate whenever K is too large. On the other hand, stability cannot in general detect when K is too small. In order to overcome this limitation, we introduce the concept of *within-cluster stability*.

4 Between-cluster and within-cluster stability

A clustering algorithm applied with the same parameters to perturbed versions of a data set should find the same structure and obtain similar results. The stability principle described by (1) relies on between-cluster boundaries and we thus call it *between-cluster stability*. Therefore, it cannot detect structure within clusters. In Figure 1, $K = 2$ is stable, whereas one cluster contains two sub-clusters. This sub-structure cannot be detected by between-cluster stability alone. Obviously, this implies that stability is unable to decide whether a data set is clusterable

or not (i.e. when $K^* = 1$), which is a severe limitation. For this very reason, we introduce a second principle of *within-cluster stability*: clusters should not be composed of several sub-clusters. This implies the absence of stable structures inside any cluster. In other words, any partition of a cluster should be unstable. The combination of these two principles leads to a new definition of a clustering:

Definition 1. *A clustering is a partitioning of data into groups so that the partition is stable, and within each cluster, there exists no stable partition.*

A clustering should have a high between-cluster stability and a low within-cluster stability. Despite its apparent simplicity, implementing this principle is a difficult task. As seen in the last section, between-cluster stability can be estimated in many ways. On the other hand, within-cluster stability is a challenging quantity to define and estimate. We propose a method to estimate both quantities, and then we detail and discuss our choices.

4.1 Stadion: a novel stability-based validity index

Let $\{\mathbf{X}_1, \dots, \mathbf{X}_D\}$ be D perturbed versions of the data set obtained by adding random noise to the original data set \mathbf{X} . Between-cluster stability of algorithm \mathcal{A} with parameter K estimates the expectation (1) by the empirical mean of the similarities s between the reference clustering $\mathcal{C}_K = \mathcal{A}(\mathbf{X}, K)$ and the clusterings of the perturbed data sets,

$$\text{Stab}_B(\mathcal{A}, \mathbf{X}, K) := \frac{1}{D} \sum_{d=1}^D s(\mathcal{A}(\mathbf{X}, K), \mathcal{A}(\mathbf{X}_d, K)). \quad (2)$$

Since s is a similarity measure, this quantity needs to be maximized. In order to define within-cluster stability, we need to assess the presence of stable structures inside each cluster. To this aim, we propose to *cluster again* the data within each cluster of \mathcal{C}_K . Formally, let $\Omega \subset \mathbb{N}^*$ be a set of numbers of clusters. The k -th cluster in the reference clustering is noted C_k , its number of elements N_k . Within-cluster stability of algorithm \mathcal{A} is defined as

$$\text{Stab}_W(\mathcal{A}, \mathbf{X}, K, \Omega) := \sum_{k=1}^K \left(\frac{1}{|\Omega|} \sum_{K' \in \Omega} \text{Stab}_B(\mathcal{A}, C_k, K') \right) \times \frac{N_k}{N}. \quad (3)$$

As a good clustering is unstable within each cluster, this quantity needs to be minimized. Hence, we propose to build a new validity index combining between-cluster and within-cluster stability. A natural choice is to maximize the difference between both quantities. We call this index *Stadion*, standing for *stability difference criterion*:

$$\text{Stadion}(\mathcal{A}, \mathbf{X}, K, \Omega) := \text{Stab}_B(\mathcal{A}, \mathbf{X}, K) - \text{Stab}_W(\mathcal{A}, \mathbf{X}, K, \Omega). \quad (4)$$

The same partition $\mathcal{C}_K = \mathcal{A}(\mathbf{X}, K)$ is used in both terms of (4). Thus, Stadion evaluates the stability of an algorithm w.r.t. a reference partition.

How to perturb data? We consider the setting in Figure 2 that is deemed to be realistic. Neither jumping nor jittering will occur if data are perturbed by sampling processes, as soon as there is enough data. Therefore, only noise-based perturbation is considered here. Among them, we adopt the ε -Additive Perturbation (ε -AP) with Gaussian or uniform noise, assuming variables are scaled to zero mean and unit variance. The number of perturbations D can be kept very low and still gives reliable estimates (an analysis on the influence of D is conducted in the supplementary material).

How to choose ε ? A central trade-off has to be taken into account when perturbing the data set. If the noise level ε is too strong, we might alter the very structure of the data. We propose to circumvent this issue by *not* choosing a single value for ε , but a grid of values. By gradually increasing ε from 0 to a value ε_{\max} , we obtain what we call a *stability path*, i.e. the evolution of stability as a function of ε . This method has one crucial advantage: it allows to compare partitions for different values of ε without the necessity of choosing one. However, it comes with two drawbacks: setting both the fineness and the maximum value of the grid. In our experiments, the fineness does not play a major role in the results. A straightforward method to fix a maximum value ε_{\max} beyond which comparisons are not meaningful anymore is as follows. The perturbation corresponding to ε_{\max} is meant to destroy the cluster structure of the original data. This corresponds to the value where the data are no longer clusterable, i.e. $K = 1$ becomes the solution with the best Stadion value. A first guess at $\varepsilon_{\max} = \sqrt{p}$ (where p is the data dimension) works well in practice. We found that visualizing the stability paths (see Figure 3) greatly helps interpreting the structures found by an algorithm, hence improving the *usefulness* of results.

How to compare partitions? The similarity measure s chosen to compare two partitions is the ARI. Note that it is used to compare cluster assignments and not the ground-truth labels. Its value is in $[0, 1]$, thus Stadion has a value in $[-1, 1]$, with 1 corresponding to the best clustering and -1 to the worst. A total of 16 different similarity measures (such as the NMI) were compared (results of this study are in supplementary material).

How to aggregate the Stadion path? To compute a scalar validity index for model selection, the Stadion path must be aggregated on the noise strength ε from 0 to ε_{\max} . Two aggregation strategies, the maximum (Stadion-max) and the mean (Stadion-mean), are evaluated in our experiments.

The within-cluster stability is governed by the parameter Ω , which detects stable structures inside clusters of \mathcal{C}_K . As these are unknown, averaging several different values in Ω gives a better estimate. In absence of sub-clusters, all partitions will be unstable because cluster boundaries will be placed in high-density regions. For the opposite reason, in presence of sub-clusters, at least some partitions will result in higher stability, thus increasing the within-cluster stability. The analysis of influence conducted in supplementary material shows that Ω has low impact on Stadion results and can be set easily.

An important assumption behind our implementation of within-cluster stability is that for non-clusterable structures (e.g. uniform noise), the algorithm will place cluster boundaries in high-density regions to produce instability through jittering. This encompasses a wide range of algorithms such as center-based, spectral or Ward linkage clustering which, for the sake of saving cost, would cut through dense clouds of points. If this requirement is not fulfilled, further studies are needed to determine whether this method will work.

5 Experiments

5.1 A simple example with stability paths

We begin by illustrating our method with K -means and uniform ε -AP on the example discussed previously (see Figure 1). Figure 3 displays between-cluster stability, within-cluster stability and Stadion as a function of the noise strength ε . For reasonable amounts of noise, the solutions $K = 1$, $K = 2$ and $K = 3$ are all perfectly stable, showing the insufficiency of between-cluster stability alone to indicate whenever K is too small. The solutions for $K \geq 4$ cut through the clusters and are thus unstable due to jittering. However, the solutions for $K = 1$ and $K = 2$ both have high within-cluster stability, caused by the presence of sub-clusters, which is not the case for $K \geq 3$. By computing a difference, our criterion Stadion combines this information and is able to indicate the correct number of clusters ($K = 3$) by selecting the Stadion path with the highest maximum or mean value. Through its formulation, Stadion is acting as a *stability trade-off*. The stability paths also give additional insights about the data structure. For example, we can read from the between-cluster stability path how the clusters successively merge together as ε increases. Finally, the last graph (called stability trade-off plot) represents Stadion-mean for different values of K .

5.2 Benchmark of clustering validation methods

Methodology. Importantly, we aim at evaluating clustering validation methods and not the clustering algorithms themselves. Thus, we evaluate methods on a large collection of 73 artificial benchmark data sets, most of them extensively used in literature, with a guaranteed known ground-truth cluster structure. Most data sets are available in [5,22] and all data will be shared after publication. It was ensured that the evaluated algorithms are able to obtain good solutions (i.e., reasonably close to the ground-truth clustering), for some optimal parameter K . The data sets also provide various difficulty levels by varying the numbers, sizes, variances, shapes of clusters and noisy, close-by or overlapping clusters.

To compare the different validation methods, we first report the number of data sets where each method found K^* , which we refer to as the number of *wins*. However, only checking whether K^* is selected is not always related to the goodness of the partition w.r.t. ground-truth, as the algorithm does not necessarily succeed in finding a good partition into K^* clusters.




Fig. 3. Between-cluster stability paths (top left), within-cluster stability paths (top right), Stadion paths (bottom left) and stability trade-off curve (bottom right) for K -means on the data set shown in Figure 1, for $K \in \{1 \dots 6\}$. ε is the amplitude of the uniform noise perturbation. The best solution $K = 3$ is selected either by taking the maximum or by averaging Stadion over ε until ε_{\max} . The trade-off plot represents the averaged Stadion, between- and within-cluster stability as a function of K .

Thus, we also compute the ARI between the selected partition and the ground-truth. Let us note $\mathcal{Y}_{K^*} = \{Y_1, \dots, Y_{K^*}\}$ the ground-truth partition. The performance of each method is assessed by computing $\text{ARI}(\mathcal{Y}_{K^*}, \mathcal{C}_{\hat{K}})$, where \hat{K} is the estimated number of clusters. In order to compare methods over multiple data sets, we compute the average ranks, denoted $\overline{R}_{\text{ARI}}$. Since data sets have different difficulties, their results are not comparable and simply reporting an average would be meaningless [16]. Thus, comparing their ranks is a more sound and fair solution.

In this benchmark, three algorithms are considered: K -means, Gaussian Mixture Models (GMM) and Ward hierarchical clustering. For K -means, two versions of Stadion are evaluated: the first one using the stability computation described in Section 4.1 (referred to as the *standard* version), and the second one with an approximation using the extension operator (referred to as the *extended* version). As seen in Section 3, an extension operator extends a clustering to new data points. K -means extends naturally by computing the Euclidean distance to centers. Hence, instead of re-running K -means for each perturbation, we directly predict the cluster assignments of perturbed data points. GMM allows a

Table 1. Benchmark results on 73 data sets for K -means, Ward and GMM. Average rank of the ARI between the selected clustering and ground-truth clustering ($\overline{R}_{\text{ARI}}$) and number of times the ground-truth K^* was selected (wins).

Method	K -means		Ward		GMM	
	$\overline{R}_{\text{ARI}}$	wins	$\overline{R}_{\text{ARI}}$	wins	$\overline{R}_{\text{ARI}}$	wins
K^* (Oracle)	8.11	73	4.77	73	5.05	73
Stadion-max	7.46	50	5.25	54	-	-
Stadion-mean	7.70	51	5.80	49	-	-
Stadion-max (extended)	7.58	56	-	-	5.59	56
Stadion-mean (extended)	8.09	48	-	-	6.79	43
BIC	-	-	-	-	6.45	48
Wemmert-Gancarski [17]	8.33	53	5.40	54	5.77	52
Silhouette [40]	9.55	46	6.47	45	7.01	45
Lange [32]	10.18	45	6.53	51	6.99	48
Davies-Bouldin [15]	10.21	40	6.45	41	7.29	34
Ray-Turi [39]	10.28	37	6.97	40	7.68	33
Hennig [26]	10.72	37	-	-	-	-
Calinski-Harabasz [14]	11.44	41	7.14	39	7.43	37
Gap statistic (B) [46]	11.49	29	-	-	-	-
X-means [38]	11.56	28	-	-	-	-
Dunn [18]	13.09	26	7.77	33	7.92	34
Hofmeyr [28]	13.20	30	-	-	-	-
Xie-Beni	13.30	22	7.61	34	8.19	28
Gap statistic (A) [46]	13.57	26	-	-	-	-
G-means [23]	13.74	24	-	-	-	-
Ben-Hur [11]	14.34	20	7.86	31	8.85	28
SpecialK [27]	17.07	19	-	-	-	-

similar extension, by assigning points to the cluster with the highest posterior probability. It is the only version considered due to GMM’s computational cost.

Table 1 summarizes results for each algorithm and validation method. We evaluated $K \in \{1, \dots, 60\}$. For Stadion, we used uniform noise, $D = 10$, $\Omega = \{2, \dots, 10\}$ and $s = \text{ARI}$. We also evaluate the partitions obtained with the ground-truth K^* and a selection of widely used clustering validation indices (when applicable) [17], the Gap statistic [46] (with alternative versions A and B implemented in [34]), BIC, X-means [38], G-means [23], the Hennig procedure [26], stability methods [11,32], and the recent SpecialK [27]. For SpecialK, we used the default parameters indicated by the authors, but the assumptions made by the method failed on 11 data sets, explaining the poor results. Unfortunately, other methods like dip-means or skinny-dip did not have easy-to-use available implementations and were not included in this study.

Stadion-max achieves the best results overall. On K -means, it is even ranked higher than the Oracle in terms of ARI. The second-best performing index is Wemmert-Gancarski (WG). It was shown in [3] that agglomerative clustering is not robust to noise, which explains inferior Stadion results with Ward. Moreover,

results are slightly biased in favor of the indices that are only valid for $K \geq 2$, unlike Stadion that will select $K = 1$ on non-clusterable distributions.

6 Conclusion

In this paper, we tackled some of limitations of cluster stability for model selection. Our contribution is twofold. First, stability can be well estimated through additive noise perturbation, solving the limitations of sampling-based stability methods. Second, we introduce the Stadion (stability difference criterion), a novel criterion acting as a trade-off between traditional stability and within-cluster stability. Furthermore, our method to control the amount of perturbation provides an interpretable visualization called stability paths.

We evaluated Stadion and methods of the literature on 73 clustering benchmark data sets. Performance is superior or on par with internal clustering indices that were designed with specific cluster geometries in mind, while relying on more general assumptions. This comes at a computational cost, requiring to run the algorithm many times. Nevertheless, studies have shown that it can be drastically reduced by down-sizing the hyperparameters with negligible impact on performance. Moreover, most theoretical results used here were derived for K -means, and more work is needed to extend these concepts to other algorithms.

Altogether, model selection remains a challenge and there is yet no theory nor a methodology that can fulfill this task perfectly. We proposed an empirical method showing interesting results along with hints to a theoretical background that could be established in future work. We hope that it will spark much-needed interest in the research community to further advance this field.

References

1. Arel-Bundock, V.: xclara dataset. <http://vincentarelbundock.github.io/Rdatasets/datasets.html>
2. Arthur, D., Vassilvitskii, S.: k-means++: The Advantages of Careful Seeding. In: ACM-SIAM Symposium on Discrete Algorithms (2007)
3. Balcan, M.F., Liang, Y.: Clustering under perturbation resilience. *SIAM Journal on Computing* (2016)
4. Bandyopadhyay, S., Pal, S.K.: Classification and Learning Using Genetic Algorithms: Applications in Bioinformatics and Web Intelligence. Springer, Heidelberg (2007)
5. Barton, T.: <https://github.com/deric/clustering-benchmark>
6. Ben-David, S.: Clustering-what both theoreticians and practitioners are doing wrong. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
7. Ben-David, S., Pál, D., Simon, H.U.: Stability of k-means clustering. In: International conference on computational learning theory (2007)
8. Ben-David, S., Von Luxburg, U.: Relating clustering stability to properties of cluster boundaries. 21st Annual Conference on Learning Theory, COLT 2008 (2008)
9. Ben-David, S., Von Luxburg, U., Pál, D.: A sober look at clustering stability. In: International Conference on Computational Learning Theory (2006)

10. Ben-David, S., Reyzin, L.: Data stability in clustering: A closer look. *Theoretical Computer Science* (2014)
11. Ben-Hur, A., Elisseeff, A., Guyon, I.: A stability based method for discovering structure in clustered data. *Pacific Symposium on Biocomputing* (2002)
12. Benavoli, A., Corani, G., Mangili, F.: Should we really use post-hoc tests based on mean-ranks? *The Journal of Machine Learning Research* (2016)
13. Bubeck, S., Meila, M., Luxburg, U.V.: How the initialization affects the stability of the k-means algorithm. *ESAIM - Probability and Statistics* (2012)
14. Caliński, T., Harabasz, J.: A dendrite method for cluster analysis. *Communications in Statistics* (1974)
15. Davies, D.L., Bouldin, D.W.: A Cluster Separation Measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1979)
16. Demšar, J.: Statistical comparisons of classifiers over multiple data sets. *Journal of machine learning research* (2006)
17. Desgraupes, B.: ClusterCrit: Clustering Indices. CRAN Package (2013)
18. Dunn, J.C.: Well-Separated Clusters and Optimal Fuzzy Partitions. *Journal of Cybernetics* (1974)
19. Falasconi, M., Gutierrez, A., Pardo, M., Sberveglieri, G., Marco, S.: A stability based validity method for fuzzy clustering. *Pattern Recognition* (2010)
20. Fang, Y., Wang, J.: Selection of the number of clusters via the bootstrap method. *Computational Statistics and Data Analysis* **56**(3), 468–477 (2012)
21. Friedman, M.: A comparison of alternative tests of significance for the problem of m rankings. *The Annals of Mathematical Statistics* (1940)
22. Gagolewski M., Bartoszuk M., C.A.G.: A new, fast, and outlier-resistant hierarchical clustering algorithm (2016)
23. Hamerly, G., Elkan, C.: Learning the K in K-means. *NIPS* (2004)
24. Handl, J.: Cluster generators. <https://personalpages.manchester.ac.uk/staff/Julia.Handl/generators.html> (2004)
25. Handl, J., Knowles, J.: Multiobjective clustering with automatic determination of the number of clusters (2004)
26. Hennig, C.: Cluster-wise assessment of cluster stability. *Computational Statistics & Data Analysis* **52**(1), 258–271 (2007)
27. Hess, S., Duivesteijn, W.: K Is the Magic Number - Inferring the Number of Clusters Through Nonparametric Concentration Inequalities. In: *EMCL-PKDD* (2019)
28. Hofmeyr, D.P.: Degrees of freedom and model selection for k-means clustering. *arXiv preprint arXiv:1806.02034* (2018)
29. Holm, S.: A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics* (1979)
30. Hoos, H., Leyton-Brown, K.: An efficient approach for assessing hyperparameter importance. In: *International Conference on Machine Learning* (2014)
31. Katti Faceli, Tiemi C. Sakata, M.C.P.d.S., de Carvalho, A.C.P.L.F.: Partitions selection strategy for set of clustering solutions. *Neurocomputing* (2010)
32. Lange, T., Roth, V., Braun, M.L., Buhmann, J.M.: Stability-based validation of clustering solutions. *Neural Computation* (2004)
33. M. C. Su, C.H.C., Hsieh, C.C.: Fuzzy C-Means Algorithm with a Point Symmetry Distance. *International Journal of Fuzzy Systems* (2005)
34. Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., et al.: Package ‘cluster’ (2013)
35. Meila, M.: How to tell when a clustering is (approximately) correct using convex relaxations. In: *Advances in Neural Information Processing Systems* (2018)
36. Möller, U., Radke, D.: A cluster validity approach based on nearest-neighbor resampling. *Proceedings - International Conference on Pattern Recognition* (2006)

37. P. Fränti, O.V.: Iterative shrinking method for clustering problems. *Pattern Recognition* (2006)
38. Pelleg, D., Moore, A.: X-means: Extending K-means with Efficient Estimation of the Number of Clusters. In: *International Conference on Machine Learning* (2000)
39. Ray, S., Turi, R.: Determination of number of clusters in k-means clustering and application in colour image segmentation. *Proceedings of the 4th international conference on advances in pattern recognition and digital techniques* (1999)
40. Rousseeuw, P.J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* (1987)
41. Salvador, S., Chan, P.: Determining the Number of Clusters/Segments in Hierarchical clustering/Segmentation Algorithm. In: *ICTAI*
42. Scrucca, L., Raftery, A.E.: Improved initialisation of model-based clustering using gaussian hierarchical partitions. *Advances in data analysis and classification* (2015)
43. Shamir, O., Tishby, N.: Cluster stability for finite samples. In: *Advances in Neural Information Processing Systems* (2007)
44. Smith, S.P., Dubes, R.: Stability of a hierarchical clustering. *Pattern Recognition* (1980)
45. Strauss, J.S., Bartko, J.J., Carpenter, W.T.: The use of clustering techniques for the classification of psychiatric patients. *British Journal of Psychiatry* (1973)
46. Tibshirani, R., Walther, G., Hastie, T.: Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society Series B* (2001)
47. Ultsch, A.: Clustering with SOM: U*C. In: *Workshop on Self Organizing Feature Maps* (2005)
48. Veenman, C.J., M.R., Backer, E.: A maximum variance cluster algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2002)
49. Vijayaraghavan, A., Dutta, A., Wang, A.: Clustering stable instances of euclidean k-means. In: *Advances in Neural Information Processing Systems* (2017)
50. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *The Journal of Machine Learning Research* (2010)
51. Von Luxburg, U.: Clustering stability: An overview. *Foundations and Trends® in Machine Learning* (2010)
52. Wilcoxon, F.: Probability tables for individual comparisons by ranking methods. *Biometrics* (1947)
53. Yeung, K.Y., Haynor, D.R., Ruzzo, W.L.: Validating clustering for gene expression data. *Bioinformatics* (2001)
54. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: *Advances in Neural Information Processing Systems* (2004)
55. Zhao, Q., Xu, M., Fränti, P.: Extending external validity measures for determining the number of clusters. *Intelligent Systems Design and Applications* (2011)



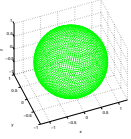
Selecting the Number of Clusters K with a Stability Trade-off: an Internal Validation Criterion – Supplementary Material

1 Additional experiments and examples

1.1 Finding $K = 1$: the case of non-clusterable data

Is a data set clusterable? Between-cluster stability is unable to answer this question, as the solution with a single cluster is trivially stable. Some stability methods are not even defined for $K = 1$ because of normalization [32]. Moreover, many internal indices use between-cluster distances and are not defined for a single cluster neither. We verified empirically that our criterion consistently outputs $K = 1$ in the case when the algorithm does not find any cluster structure. Table 1 contains results for non-clusterable distributions. Stadion outputs $K = 1$ in all cases. An example of Stadion path and trade-off curve for the golfball data set is provided in Figure 1 (results are similar for other data sets).

Table 1. Number of clusters found by Stadion on non-clusterable distributions.

Dataset		N	dimension	K selected by Stadion
Uniform cube (2d)		1000	2	1
Uniform cube (10d)		1000	10	1
Gaussian (2d)		1000	2	1
Gaussian (10d)		1000	10	1
Golfball [47]		4002	3	1




Fig. 1. Stadion path (left) and stability trade-off plot (right) on the golfball data set with K -means. $K = 1$ is clearly selected by Stadion-max/mean (uniform noise, $\Omega = \{2, \dots, 10\}$).

1.2 Examples of jumping between local minima

As explained in Section 3, two sources of instability are jumping and jittering. We have already stated that our method leverages jittering of cluster boundaries in high-density regions due to perturbation. Jumping, on the other hand, happens when the algorithm finds very different solutions on different samples; in case of objective-minimizing algorithms, it ends up in different local minima. Two main effects lead to jumping: initialization, and symmetries in the data distribution. Finally, subtle geometrical properties of the distribution might also cause jumping [51]. An example of jumping of K -means due to symmetries is shown on




Fig. 2. Example of K -means jumping between three global minima for $K = 2$ on a symmetric distribution with three Gaussians, despite consistent initialization (K -means++ and best of 10 runs). Under slight perturbation (here uniform ϵ -AP, but resampling gives identical results), the algorithm jumps between grouping two random clusters together.

Figure 2: clearly, there are several global minima, and even if the algorithm is deterministic, slight perturbations of the distribution (noise or sampling) make

the algorithm jump between solutions. The second cause of jumping is due to initialization. As illustrated by Figure 3 for K -means, if a single random initialization is used, depending on the initial position of centers, four different configurations occur randomly, even without any perturbation of the data. We




Fig. 3. Example of K -means jumping between three local minima for $K = 4$, when a single random initialization is used. Depending on the initial centers configuration, the algorithm jumps between splitting a random cluster in two (a, b, c).

place ourselves in a realistic setting without perfect symmetries and an effective algorithm initialization strategy, thus jumping is not the main source of instability.

1.3 Failure of sampling-based stability methods

In this section, we will see on a trivial example why stability methods based on re-sampling are not reliable to detect the presence of structure in the data. Four methods are compared:

1. Stadion based on ε -Additive Perturbation
2. Stadion based on bootstrapping perturbation
3. The model explorer algorithm [11] based on subsampling
4. The model order selection method [32] based on splitting data in two halves and transferring labels from one half onto the other using a supervised nearest-neighbor classifier.

We demonstrate that only the first method is successful on a simple example consisting in a mixture of two correlated Gaussians, represented on Figure 4. Data are scaled to zero mean and unit variance as for every other data set. K -means is used to cluster the data. As illustrated on the plot, K -means with $K = 2$ separates almost perfectly the two Gaussians. All other solutions split the two Gaussians into several sub-clusters of equal sizes, with cluster boundaries lying in the regions of highest density, as can be seen from the example for $K = 4$ (where the boundaries are in the middle of the Gaussians). Thus, in addition to being the best solution, $K = 2$ is the only acceptable one. However, sampling-based methods fail in assessing its stability, since they estimate $K = 4$ as the

most stable solution. This result can be explained because the data set is not symmetric and for each K there is one global minimum so no jumping occurs, even with a random initialization scheme. Thus the only possible source of instability stems from jittering. As expected in theory, our experiments show that the different sampling processes did not succeed in creating jittering. Conversely, ε -AP has indeed produced jittering. In details, the model order selection method




Fig. 4. Example data set of two correlated Gaussians, scaled to zero mean and unit variance. With the K -means algorithm, all sampling-based methods select $K = 4$ or $K = 6$, whereas with ε -Additive Perturbation, $K = 2$ is the only stable solution.

[32] selects $K = 4$, followed by $K = 6$. The model explorer [11] finds $K = 6$ as the best solution, followed by $K = 4$. These results are consistent across initialization schemes (random, K -means++, best of several runs). Hence, random initialization will not help creating instability by jumping. Stadion with additive noise was able to find $K = 2$ among the set of tested values $\{1, \dots, 6\}$ (see Figure 5). This is not only due to adding the within-cluster stability. As evidence, we replaced ε -AP by a bootstrap perturbation: Stadion with bootstrapping also fails, selecting $K = 1$ as the best solution followed by $K = 4$, and this for all initialization schemes.




Fig. 5. Between-cluster stability, within-cluster stability and Stadion paths (uniform noise, $\Omega = \{2, \dots, 6\}$) on the example of two correlated Gaussians where all sampling-based methods fail. Stadion clearly finds $K = 2$ by taking the max or mean of the path curve.

1.4 Example of Stadion behavior with K -means

This example illustrates the behavior of our stability criterion Stadion and how to interpret the stability paths, using the data set 2d-4c shown in Figure 6. It consists in four clusters with different variance and size, where two clusters are closer to each other while the other clusters are at a greater distance. At first glance, this example looks trivial, but the majority of internal indices fail. For instance, the Dunn and Silhouette indices both select $K = 3$. The stability




Fig. 6. Between-cluster stability, within-cluster stability and Stadion paths (uniform noise, $\Omega = \{2, \dots, 6\}$) on the 2d-4c data set. Stadion selects $K = 4$, followed by $K = 3$.

paths in Figure 6 show that Stadion is able to detect the structure of the data and selects $K = 4$. The only difference between the solutions with $K = 4$ and $K = 5$ is that the largest cluster (in green) is split, thus leading to a much lower between-cluster stability but the same within-cluster stability. Solutions $K = 2$ and $K = 3$ group clusters together without any splitting. Therefore, those solutions have a high between-cluster stability and also a high within-cluster stability. Altogether, on the Stadion path (Figure ??), the path corresponding to $K = 4$ is similar to $K = 5$ whereas $K = 2$ and $K = 3$ have an equivalent behavior. This is due to the structure of the data, and especially because the two rightmost clusters are close to each other. The moment when the path of

solution $K = 3$ becomes the best solution is the moment when these two clusters merge because of a high ε -AP, and this is also the moment where $K = 1$ prevails.

Finally, Stadion paths (with stability and instability paths) give useful additional information on a clustering and on the structure of the data. When $K > K^*$, the paths are similar to the path of K^* but with a smaller scale, as they have the same within-cluster stability but lower between-cluster stability. On the other hand, when $K < K^*$, the paths are shifted towards the right, and may become superior for larger ε values.

1.5 Whenever K^* is not the best partition

Sometimes, the best solution is not the one obtained with the true parameter K^* , because the algorithm itself is unable to recover the ground-truth partition. This is the case for the `4clusters_corner` data set, depicted on Figure 7. While the most natural solution is to separate the four clusters, it is not achievable by K -means: with $K^* = 4$, it will cut through the large cluster instead of separating the two small green clusters, for the sake of saving the cost induced by the variance and the size of this cluster. Among the proposed solutions, the highest ARI (w.r.t. the ground-truth) is obtained with $K = 3$ (ARI = 0.92), followed by $K = 2$ (0.74), $K = 5$ (0.65) and lastly $K^* = 4$ (0.58). Almost all internal indices select





Fig. 7. Solutions of K -means on the `4clusters_corner` data set for $K \in \{2, \dots, 5\}$.

$K = 2$. Stability methods based on sampling [11,32] selected the ground-truth $K^* = 4$, earning them a "win", although it is the worst partition among the four. We explain it by the fact that these methods are not leveraging jittering inside the large cluster. Finally, Stadion always selects the solution $K = 3$ having the highest ARI. Moreover, the criterion outputs solutions in the same order than ARI. This examples clearly exhibits the stability trade-off occurring in Stadion: it tries to preserve a high between-cluster stability while keeping within-cluster stability as low as possible (see Table 2). Stadion paths on Figure 8 also show how the three smaller clusters merge as the noise level increases.

Table 2. Stability trade-off leveraged by Stadion on the 4clusters_corner data set.

K	ARI	StabB	StabW	Stadion
1	0.00	++	--	0 ✗
2	0.74	++	-	+ ✗
3	0.92	++	+	+++ ✓
4	0.58	--	+	- ✗
5	0.65	--	++	0 ✗

**Fig. 8.** Stadion paths on the 4clusters_corner data. $K = 3$ is selected although $K^* = 4$.

1.6 Example of K approaching N

This paragraph introduces the behavior of Stadion when the number of clusters K evaluated becomes as large as the number of samples N . Even if this is beyond the common setting in clustering, the criterion is still consistent. Figure 9 displays the stability trade-off for K -means on an example with three Gaussians, using ARI as the similarity metric for stability estimation. As K approaches N ,

1. Between-cluster stability decreases towards 0, except for $K = N$ where it jumps back to 1, because all partitions with one sample per cluster are perfectly similar to ARI.
2. Within-cluster stability increases towards 1, as clusters with few samples become trivially stable.
3. Stadion still indicates the correct solution $K = 3$, while decreasing towards -1 , only jumping back to 0 when $K = N$.

Note that the borderline case $K = N$ depends on the similarity measure s used.




Fig. 9. Stability trade-off plot for K -means on three Gaussians with $N = 50$ for $K \in \{1, \dots, 50\}$ (uniform noise, $\Omega = \{2, \dots, 10\}$). Stadion is still valid when the tested K becomes large.

2 Hyperparameter study

The stability difference criterion (Stadion) introduced in this work is governed by several hyperparameters:

- D : the number of perturbed samples used in the stability computation.
- noise: the type of noise for ε -Additive Perturbation. We experimented with uniform and Gaussian noise.
- Ω : the set of algorithm parameters K' used in within-cluster stability.

The goal of this section is to study their importance and impact on the performance of Stadion for clustering model selection, using the three studied algorithms (K -means, Ward linkage and GMM). Only the extended versions of Stadion for K -means and GMM are included to speed up the analysis. Performance is evaluated using the ARI between the selected solution and the ground-truth clustering.

2.1 Importance study with fANOVA

Ideally, practitioners would like to know how hyperparameters affect performance, not just in the context of a single fixed instantiation of the remaining hyperparameters, but across all their instantiations. The fANOVA (functional ANalysis Of VAriance) framework for assessing hyperparameter importance introduced in [30] is based on efficient marginalization over dimensions using regression trees. The importance of each hyperparameter is obtained by training a Random Forest model of 100 regression trees to predict the performance of Stadion given the set of hyperparameters. Then, the variance of the performance due to a given hyperparameter is decomposed by marginalizing out the effects of all

other parameters. It also allows to assess interaction effects. Hence, the fANOVA framework provides insights on the overall importance of hyperparameters and their interactions.

The maximum amount of noise ε_{\max} and the fineness of the grid are not included in the study, because it is data-dependent and one can easily check if values are appropriate by looking at the paths. We study the following discrete hyperparameter space:

- $D \in \{1, \dots, 10\}$
- noise is uniform or Gaussian
- $\Omega \in \{2, 3, 5, 10, \{2, \dots, 5\}, \{2, \dots, 10\}, \{10, \dots, 20\}, \{2, \dots, 20\}\}$




Fig. 10. Box plots of the fANOVA importance of parameters and their interactions for Stadion-max (left) and mean (right) across 73 benchmark data sets, for three algorithms.

Figure 10 shows the contributions of hyperparameters and their interactions to the variance of ARI performance, across 73 benchmark data sets. Ω is by far the most important parameter, and this for two reasons.

First, the size of the data N needed to obtain good estimations is relative to the number of clusters K . More precisely, in our theoretical setting we consider $K \ll N$. Whenever K' is large relative to N_k , which can happen in within-cluster stability for small clusters, jumping can occur. This implies that even in presence of sub-clusters, high values of K' in Ω will create instability and thus lead to low within-cluster stability. More precisely, if $K \geq K^*$, then in general Ω will not affect within-cluster stability because it is already low. But whenever $K \leq K^*$, within-cluster stability is more impacted by large values of K' in Ω , negatively affecting Stadion. However, impact on performance is very limited as shown in the following paragraph.

The second most important parameter is the interaction (D, Ω) , for the same reason: large numbers of clusters make estimating the within-cluster stability more difficult, and thus a higher number of perturbations D is needed to obtain a good approximation.

2.2 Influence of D

The D hyperparameter defines the number of perturbed samples used in the stability computation. In our benchmark, we used $D = 10$. Surprisingly, a number of perturbed samples as low as $D = 1$ already gives a good estimate of the expectation and the performance only slightly increases with larger values of D . We perform an experiment by making D vary from 1 to 10, keeping other hyperparameters fixed (uniform noise, $\Omega = \{2, \dots, 10\}$), for the three algorithms and both Stadion path aggregation strategies (max and mean), and measure performance in terms of ARI over 73 benchmark data sets. Results on Figure 11 show that low D values have a higher variance and slightly lower performance. To quantify further the influence of this parameter, we followed the recommen-

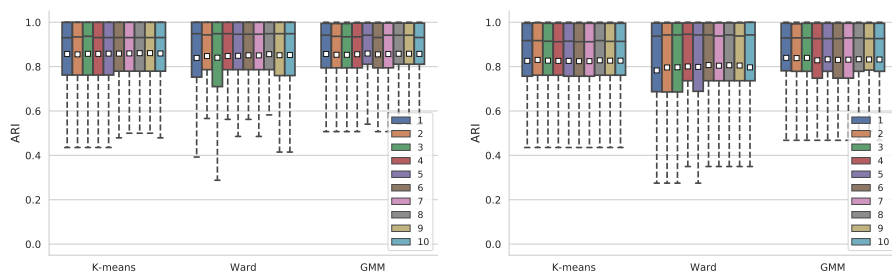


Fig. 11. Box-plot of the ARI of partitions selected by Stadion-max (left) and mean (right) across 73 data sets, for three algorithms and different values of D , the number of samples in the stability computation.

dation in [16] and used the Friedman test [21] for comparisons on multiple data sets, in order to test against the null hypothesis \mathcal{H}_0 stating that all parameters have equivalent performance. After rejecting \mathcal{H}_0 , we performed the pairwise post-hoc analysis recommended by [12] where the average rank comparison (e.g. Nemenyi test) is replaced by a Wilcoxon signed-rank test [52] at $\alpha = 5\%$ with a Holm-Bonferroni correction procedure to control the family-wise error rate (FWER) [29,?]. To visualize post-hoc test results, we use the critical difference (CD) diagram [16], where a thick horizontal line shows groups (cliques) of classifiers that are not significantly different in terms of performance. In all but one case, the Friedman test could not reject the null hypothesis. Only for the GMM algorithm and max aggregation, the null hypothesis was rejected, leading to the critical difference diagrams on Figure 12. The number of samples D has a negligible impact on the performance of our method. We assume it is due to the fact that in our setting, instability is caused by jittering at cluster boundaries, which does not vary much from one perturbation to another with reasonable amounts of data. On the contrary, sampling-based stability methods that rely on jumping require a much higher number of samples (for instance, [11] use 100 samples




Fig. 12. Critical difference diagrams after Wilcoxon-Holms test ($\alpha = 5\%$) on GMM performance, for Stadion-max with uniform (left) and Gaussian (right) noise, for different values of D , the number of perturbations in the stability computation.

and [32] use 20 samples). As a conclusion, we recommend using $D \geq 5$, but if computation time is costly, $D = 1$ can be used safely to cut down complexity.

2.3 Influence of noise type

We experiment with two types of ε -additive noise perturbation: uniform noise and Gaussian noise. As previously, we report the distributions of performance in terms of ARI across 73 data sets for both noise types on Figure 13 (with $D = 10$ and $\Omega = \{2, \dots, 10\}$). To assess the difference between both noise




Fig. 13. Box plots of the ARI of partitions selected by Stadion-max (left) and mean (right) across 73 data sets, for three algorithms, using uniform or Gaussian noise perturbation.

types, we perform the Wilcoxon signed-rank test on the performance results (at confidence level $\alpha = 5\%$). For every algorithm and Stadion path aggregation, the test did not reject the null hypothesis. Thus, either uniform or Gaussian noise can be used.

2.4 Influence of Ω

The Ω hyperparameter is a set defining the numbers of clusters used to cluster again each cluster of the original partition. We perform an experiment by varying Ω , keeping other hyperparameters fixed (uniform noise, $D = 10$), for both Stadion path aggregation strategies (max and mean), and measure performance




Fig. 14. Box plots of ARI of partitions selected by Stadion-max (left) and mean (right) across 73 data sets, for three algorithms and different sets Ω (numbers of clusters used in within-cluster stability).

in terms of ARI over 73 benchmark data sets. Results in Figure 14 demonstrate that Ω does not have, in most cases, a big impact on the performance of Stadion. This does not contradict the results of the fANOVA. Indeed, Ω has the largest variance contribution to the performance, but this variance remains small, and overall Stadion is robust for reasonable choices of the parameter, such as $\{2, \dots, 5\}$ or $\{2, \dots, 10\}$. Ward linkage is the most influenced by the choice of Ω . We guess the main reason is that agglomerative clustering algorithms are not robust to noise [3]. Critical difference diagrams after Wilcoxon-Holms test on performance are given in Figures (15, 16, 17). None of them showed significant differences, indicating that there is not enough data to conclude. However, small values in Ω seem to perform better, which confirms the previous claim that large values of K' in Ω negatively impact performance. In particular, the range $\{2, \dots, 10\}$ used in our benchmark performs well across all algorithms.




Fig. 15. Critical difference diagrams after Wilcoxon-Holms test on K -means performance, for different values of Ω .




Fig. 16. Critical difference diagrams after Wilcoxon-Holms test on Ward performance, for different values of Ω .




Fig. 17. Critical difference diagrams after Wilcoxon-Holms test on GMM performance, for different values of Ω . With Stadion-mean, the Friedman test could not reject \mathcal{H}_0 .

3 Complexity analysis

This section provides a time complexity analysis of the Stadion computation. Let $A(K, N)$ be the time complexity of the algorithm with parameter K and data of size N , assuming the data dimension is fixed. In addition, let $S(K, N)$ be the complexity of the similarity measure s , D the number of perturbations and M the number of noise levels.

Between-cluster stability The complexity is $\mathcal{O}((A(K, N) + S(K, N))DM)$.

Within-cluster stability For a given parameter K and a set of parameters $\Omega = \{2, \dots, K'\}$, the amount of operations is $\sum_{k=1}^K \sum_{k'=2}^{K'} (A(k', N_k) + S(k', N_k))DM$, which can be bounded by $\mathcal{O}(KK'(A(K', N) + S(K', N))DM)$.

In the case of K -means, we have $A(K, N) = \mathcal{O}(KNTI)$, where T is the number of iterations until convergence, and I the number of initialization runs. Then, ARI is linear: $S(K, N) = \mathcal{O}(N)$. Overall, we obtain a complexity for Stadion with K -means and ARI equal to $\mathcal{O}(KK'^2NTIDM)$.

The influence studies showed that Ω can be set to a small range, e.g. $\{2, \dots, 5\}$ or $\{2, \dots, 10\}$, and that D can be kept very low. In addition, the number of noise levels M is fixed. Thus, complexity in K' , D and M is manageable, and complexity of Stadion is mainly driven by the complexity of the algorithm itself.

In regard, internal indices relying on between-cluster and within-cluster distances have a complexity of $\mathcal{O}(N^2)$ with pairwise distance or $\mathcal{O}(KN)$ with centroid distance. Thus, the cost of having to run the algorithm several times may be smaller than a quadratic index if the algorithm is linear in N .

4 Experimental setting

In this section, we provide additional information on our experimental setting.

4.1 Algorithm initialization

- K -means: initialization is achieved using K -means++ [2] and keeping the best of 35 runs (w.r.t. the cost function).
- Ward linkage: agglomerative clustering is deterministic, no initialization strategy is needed.
- Gaussian Mixture Model: Two initialization schemes were considered. The first one uses K -means to initialize the EM algorithm. The second one uses the approach discussed in [42] based on a scaled SVD and agglomerative hierarchical clustering at the initialization step. However, there was no noticeable difference between the two initializations.

4.2 Preprocessing

Every data set was scaled to zero mean and unit variance on each dimension.

4.3 Choice of the external performance metric

Note that performance is evaluated using the external index ARI w.r.t. the ground-truth cluster assignments, while Stadion also uses $s = \text{ARI}$ as its similarity measure to estimate stability. It would be reasonable to expect this situation to introduce some kind of bias. However, we also measured the performance using many other external metrics (such as NMI), and almost no change in the final ranking was observed. Thus, results are not biased in favor of Stadion, and we kept ARI as a standard, adjusted, performance metric [50].

4.4 List of data sets

A complete list of the 73 data sets used in the benchmark is provided below, indicating the number of samples (N), dimension (p), ground-truth number of clusters (K^*) and reference. The data sets are easily downloaded from the archives [5] and [22]. The data sets without references are original and have been created for this work, in order to provide challenging model selection tasks. They will be available in the companion repository of this paper.

References

1. Arel-Bundock, V.: xclara dataset. <http://vincentarelbundock.github.io/Rdatasets/datasets.html>
2. Arthur, D., Vassilvitskii, S.: k-means++: The Advantages of Careful Seeding. In: ACM-SIAM Symposium on Discrete Algorithms (2007)
3. Balcan, M.F., Liang, Y.: Clustering under perturbation resilience. *SIAM Journal on Computing* (2016)
4. Bandyopadhyay, S., Pal, S.K.: *Classification and Learning Using Genetic Algorithms: Applications in Bioinformatics and Web Intelligence*. Springer, Heidelberg (2007)
5. Barton, T.: <https://github.com/deric/clustering-benchmark>
6. Ben-David, S.: Clustering-what both theoreticians and practitioners are doing wrong. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
7. Ben-David, S., Pál, D., Simon, H.U.: Stability of k-means clustering. In: *International conference on computational learning theory* (2007)
8. Ben-David, S., Von Luxburg, U.: Relating clustering stability to properties of cluster boundaries. *21st Annual Conference on Learning Theory, COLT 2008* (2008)
9. Ben-David, S., Von Luxburg, U., Pál, D.: A sober look at clustering stability. In: *International Conference on Computational Learning Theory* (2006)
10. Ben-David, S., Reyzin, L.: Data stability in clustering: A closer look. *Theoretical Computer Science* (2014)
11. Ben-Hur, A., Elisseeff, A., Guyon, I.: A stability based method for discovering structure in clustered data. *Pacific Symposium on Biocomputing* (2002)
12. Benavoli, A., Corani, G., Mangili, F.: Should we really use post-hoc tests based on mean-ranks? *The Journal of Machine Learning Research* (2016)
13. Bubeck, S., Meila, M., Luxburg, U.V.: How the initialization affects the stability of the k-means algorithm. *ESAIM - Probability and Statistics* (2012)
14. Caliński, T., Harabasz, J.: A dendrite method for cluster analysis. *Communications in Statistics* (1974)
15. Davies, D.L., Bouldin, D.W.: A Cluster Separation Measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1979)
16. Demšar, J.: Statistical comparisons of classifiers over multiple data sets. *Journal of machine learning research* (2006)
17. Desgraupes, B.: *ClusterCrit: Clustering Indices*. CRAN Package (2013)
18. Dunn, J.C.: Well-Separated Clusters and Optimal Fuzzy Partitions. *Journal of Cybernetics* (1974)
19. Falasconi, M., Gutierrez, A., Pardo, M., Sberveglieri, G., Marco, S.: A stability based validity method for fuzzy clustering. *Pattern Recognition* (2010)
20. Fang, Y., Wang, J.: Selection of the number of clusters via the bootstrap method. *Computational Statistics and Data Analysis* **56**(3), 468–477 (2012)
21. Friedman, M.: A comparison of alternative tests of significance for the problem of m rankings. *The Annals of Mathematical Statistics* (1940)
22. Gagolewski M., Bartoszek M., C.A.G.: A new, fast, and outlier-resistant hierarchical clustering algorithm (2016)
23. Hamerly, G., Elkan, C.: Learning the K in K-means. *NIPS* (2004)
24. Handl, J.: Cluster generators. <https://personalpages.manchester.ac.uk/staff/Julia.Handl/generators.html> (2004)
25. Handl, J., Knowles, J.: Multiobjective clustering with automatic determination of the number of clusters (2004)

26. Hennig, C.: Cluster-wise assessment of cluster stability. *Computational Statistics & Data Analysis* **52**(1), 258–271 (2007)
27. Hess, S., Duivesteijn, W.: K Is the Magic Number - Inferring the Number of Clusters Through Nonparametric Concentration Inequalities. In: EMCL-PKDD (2019)
28. Hofmeyr, D.P.: Degrees of freedom and model selection for k-means clustering. arXiv preprint arXiv:1806.02034 (2018)
29. Holm, S.: A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics* (1979)
30. Hoos, H., Leyton-Brown, K.: An efficient approach for assessing hyperparameter importance. In: International Conference on Machine Learning (2014)
31. Katti Faceli, Tiemi C. Sakata, M.C.P.d.S., de Carvalho, A.C.P.L.F.: Partitions selection strategy for set of clustering solutions. *Neurocomputing* (2010)
32. Lange, T., Roth, V., Braun, M.L., Buhmann, J.M.: Stability-based validation of clustering solutions. *Neural Computation* (2004)
33. M. C. Su, C.H.C., Hsieh, C.C.: Fuzzy C-Means Algorithm with a Point Symmetry Distance. *International Journal of Fuzzy Systems* (2005)
34. Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., et al.: Package ‘cluster’ (2013)
35. Meila, M.: How to tell when a clustering is (approximately) correct using convex relaxations. In: *Advances in Neural Information Processing Systems* (2018)
36. Möller, U., Radke, D.: A cluster validity approach based on nearest-neighbor re-sampling. *Proceedings - International Conference on Pattern Recognition* (2006)
37. P. Fränti, O.V.: Iterative shrinking method for clustering problems. *Pattern Recognition* (2006)
38. Pelleg, D., Moore, A.: X-means: Extending K-means with Efficient Estimation of the Number of Clusters. In: International Conference on Machine Learning (2000)
39. Ray, S., Turi, R.: Determination of number of clusters in k-means clustering and application in colour image segmentation. *Proceedings of the 4th international conference on advances in pattern recognition and digital techniques* (1999)
40. Rousseeuw, P.J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* (1987)
41. Salvador, S., Chan, P.: Determining the Number of Clusters/Segments in Hierarchical clustering/Segmentation Algorithm. In: ICTAI
42. Scrucca, L., Raftery, A.E.: Improved initialisation of model-based clustering using gaussian hierarchical partitions. *Advances in data analysis and classification* (2015)
43. Shamir, O., Tishby, N.: Cluster stability for finite samples. In: *Advances in Neural Information Processing Systems* (2007)
44. Smith, S.P., Dubes, R.: Stability of a hierarchical clustering. *Pattern Recognition* (1980)
45. Strauss, J.S., Bartko, J.J., Carpenter, W.T.: The use of clustering techniques for the classification of psychiatric patients. *British Journal of Psychiatry* (1973)
46. Tibshirani, R., Walther, G., Hastie, T.: Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society Series B* (2001)
47. Ultsch, A.: Clustering with SOM: U*C. In: *Workshop on Self Organizing Feature Maps* (2005)
48. Veenman, C.J., M.R., Backer, E.: A maximum variance cluster algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2002)
49. Vijayaraghavan, A., Dutta, A., Wang, A.: Clustering stable instances of euclidean k-means. In: *Advances in Neural Information Processing Systems* (2017)
50. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *The Journal of Machine Learning Research* (2010)

51. Von Luxburg, U.: Clustering stability: An overview. *Foundations and Trends® in Machine Learning* (2010)
52. Wilcoxon, F.: Probability tables for individual comparisons by ranking methods. *Biometrics* (1947)
53. Yeung, K.Y., Haynor, D.R., Ruzzo, W.L.: Validating clustering for gene expression data. *Bioinformatics* (2001)
54. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: *Advances in Neural Information Processing Systems* (2004)
55. Zhao, Q., Xu, M., Fränti, P.: Extending external validity measures for determining the number of clusters. *Intelligent Systems Design and Applications* (2011)

Table 3. List of the benchmark data sets.

Dataset	N	p	K^*	reference
2d-10c	2990	2	9	[24]
2d-3c-no123	715	2	3	[24]
2d-4c	1261	2	4	[24]
2d-4c-no4	863	2	4	[24]
2d-4c-no9	876	2	4	[24]
3clusters_elephant	700	2	3	
4clusters_corner	1575	2	4	
4clusters_twins	612	2	4	
5clusters_stars	1050	2	5	
A1	3000	2	20	[37]
A2	5250	2	35	[37]
curves1	1000	2	2	[5]
D31	3100	2	31	[48]
diamond9	3000	2	9	[41]
dim032	1024	32	16	[37]
dim064	1024	64	16	[37]
dim1024	1024	1024	16	[37]
dim128	1024	128	16	[37]
dim256	1024	256	16	[37]
dim512	1024	512	16	[37]
DS-577	577	2	3	[33]
DS-850	850	2	5	[33]
ds4c2sc8	485	2	8	[31]
elliptical_10_2	500	2	10	[4]
elly-2d10c13s	2796	2	10	[24]
engytime	4096	2	2	FCPS [47]
exemples1_3g	525	2	3	
exemples10_WellS_3g	975	2	3	
exemples2_5g	1375	2	5	
exemples3_Uvar_4g	1000	2	4	
exemples4_overlap_3g	1050	2	3	
exemples5_overlap2_3g	1550	2	3	
exemples6_quicunx_4g	2250	2	4	
exemples7_elbow_3g	788	2	3	
exemples8_Overlap_Uvar_5g	2208	2	6	
exemples9_YoD_6g	2208	2	6	
fourty	1000	2	40	[5]
g2-16	2048	16	2	[37]
g2-2	2048	2	2	[37]
g2-64	2048	64	2	[37]
hepta	212	3	7	FCPS [47]
long1	1000	2	2	[25]
long2	1000	2	2	[25]
long3	1000	2	2	[25]
longsquare	900	2	6	[25]
R15	600	2	15	[48]
s-set1	5000	2	15	[37]
s-set2	5000	2	15	[37]
s-set3	5000	2	15	[37]
s-set4	5000	2	15	[37]
sizes1	1000	2	4	[25]
sizes2	1000	2	4	[25]
sizes3	1000	2	4	[25]
sizes4	1000	2	4	[25]
sizes5	1000	2	4	[25]
spherical_4_3	400	3	4	[4]
spherical_5_2	250	2	5	[4]
spherical_6_2	300	2	6	[4]
square1	1000	2	4	[25]
square2	1000	2	4	[25]
square3	1000	2	4	[25]
square4	1000	2	4	[25]
square5	1000	2	4	[25]
st900	900	2	9	[4]
tetra	400	3	4	FCPS [47]
triangle1	1000	2	4	[25]
triangle2	1000	2	4	[25]
twenty	1000	2	20	[5]
twodiamonds	800	2	2	FCPS [47]
wingnut	1016	2	2	FCPS [47]
xclara	3000	2	3	[1]
zelnik2	303	2	3	[54]
zelnik4	622	2	5	[54]