

HAWKS: Evolving Challenging Benchmark Sets for Cluster Analysis

Cameron Shand, Richard Allmendinger, *Member, IEEE*, Julia Handl, Andrew Webb, and John Keane

Abstract—Comprehensive benchmarking of clustering algorithms is rendered difficult by two key factors: (i) the elusiveness of a unique mathematical definition of this unsupervised learning approach and (ii) dependencies between the generating models or clustering criteria adopted by some clustering algorithms and indices for internal cluster validation. Consequently, there is no consensus regarding the best practice for rigorous benchmarking, and whether this is possible at all outside the context of a given application. Here, we argue that synthetic datasets must continue to play an important role in the evaluation of clustering algorithms, but that this necessitates constructing benchmarks that appropriately cover the diverse set of properties that impact clustering algorithm performance. Through our framework, HAWKS, we demonstrate the important role evolutionary algorithms play to support flexible generation of such benchmarks, allowing simple modification and extension. We illustrate two possible uses of our framework: (i) the evolution of benchmark data consistent with a set of hand-derived properties and (ii) the generation of datasets that tease out performance differences between a given pair of algorithms. Our work has implications for the design of clustering benchmarks that sufficiently challenge a broad range of algorithms, and for furthering insight into the strengths and weaknesses of specific approaches.

Index Terms—Clustering, evolutionary computation, synthetic data, benchmarking, data generator.

I. INTRODUCTION

CLUSTER analysis is an unsupervised learning approach with the high-level aim of identifying groups (clusters) of objects that are more similar to each other than to the objects in other groups. It is a fundamental approach for knowledge discovery, with a broad range of applications from bioinformatics [1]–[3], to cybersecurity [4], to medicine [5].

Due to the unsupervised nature of clustering, the process of evaluating the quality of a partition (i.e. a given set of clusters) is not a straightforward task. Attempts to formally capture the qualities intuitively associated with pronounced cluster structure (such as compactness of individual clusters and separation between clusters), have led to the mathematical definition of a range of *internal* validation indices, which can be both complementary and conflicting [6]–[8]. Arbelaitz et al. [9] studied 30 such internal indices and concluded that the utility of such measures varied depending on the datasets

considered, highlighting the limited scope of each individual index.

External cluster validation indices are thought to address this limitation and to provide a more objective assessment of clustering performance [1]. However, they require knowledge of the ground truth for a given dataset, i.e. information about the correct cluster membership of each data point — information that is difficult to come by in realistic unsupervised learning applications. For this reason, synthetic benchmark datasets (i.e. datasets with a known generating model) play an important role in the evaluation of clustering performance. A key advantage of such data is that both the ground truth and any assumptions implicit to the generating process are accessible. This allows for both an objective assessment of clustering performance and informed reasoning about the key drivers behind the observed performance.

In principle, direct control over the generating model then allows for the provision of datasets with specific and varied properties. This facilitates the testing of performance with regards to these known characteristics; the benefit is concrete: practically translatable insight regarding the strengths and weaknesses of particular algorithms [10]. However, existing generators for synthetic clustering benchmarks have not been designed with this level of flexibility in mind — instead, they typically use a set of fixed (manually tuned) parameter bounds within their generating model [11], limiting the complexity and diversity of datasets that can be obtained and failing to fully match the range of challenges posed by real-world datasets.

Our framework, HAWKS, is designed to address this limitation through the integration of an evolutionary algorithm (EA) into the generating process. EAs lend themselves as a mechanism to directly control and adapt key aspects of a partition's generating model — here, we demonstrate that this facilitates the design of more powerful benchmarks, exhibiting a diverse range of properties. Furthermore, the innate modularity of EAs provides flexibility in choosing all key model components, including the representation of individual clusters and the set of objectives and constraints constituting the partition-level generating model, such as constraints on inter-cluster relationships. This flexibility is key to a broader utility of the framework, and our experiments illustrate HAWKS' potential in evolving benchmark sets either to meet a predefined set of criteria, or to directly maximize performance differences between pairs of algorithms.

In summary, the main contributions of this paper are as follows:

- 1) We propose an evolutionary framework for the generation of clustering benchmarks, HAWKS, that allows for

C. Shand is with the University College London, London WC1E 6BT, U.K. (e-mail: c.shand@ucl.ac.uk).

R. Allmendinger, J. Handl and J. Keane are with the University of Manchester, Manchester M15 6PB, U.K. (e-mail: richard.allmendinger@manchester.ac.uk, julia.handl@manchester.ac.uk, john.keane@manchester.ac.uk).

A. Webb is with vTime, Liverpool L8 5RN, U.K. (email: andrew@awebb.info).

flexible parameterization of its generating model.

- 2) We describe a set of measurable properties (problem features) that quantify the difficulty of a cluster structure from a range of perspectives. This set of problem features is utilized to define an instance space [12], enabling visual examination of the correlations between these properties and algorithm performance. For this purpose, we represent each benchmark set by its associated problem features, embed this representation into two dimensions, and use colour coding to highlight the top performing algorithm for each dataset.
- 3) We provide an indicative example of the use of HAWKS to generate datasets across the instance space by varying a subset of parameters. In this first optimization mode, individual problem features can be deployed as objectives and/or constraints. By varying the relative importance of each feature, a diverse collection of benchmarks can be obtained.
- 4) We present a second optimization mode for HAWKS, which aims to generate benchmarks that elicit performance differences between pairs of clustering algorithms. Analysis of the problem features and cluster structures associated with the resulting datasets allows for the identification of the relative strengths and weaknesses of each algorithm.

The remainder of this paper proceeds as follows. Section II further motivates the need for synthetic cluster generators, and positions our work relative to existing generators and the literature. Section III describes HAWKS, discussing the importance of each component in the framework as a whole, and the indicative design choices we have made, to illustrate its use. Section IV describes our experimental setup, including the set of problem features we deploy to measure the complexity of a dataset, and details of existing benchmark sets we compare against. The experimental results are presented in Section V, demonstrating HAWKS' ability to evolve diverse benchmark data, and instances tailored to challenge specific algorithms. Finally, we conclude and discuss future research directions in Section VI.

II. BACKGROUND

This section reviews the relevant background to our work. We start by positioning our work relative to similar problems in the literature, most prominently the relevance of benchmarks, and related work on the algorithm selection problem. We then review the issues of cluster analysis, and discuss the implications for the development and use of synthetic benchmarks. Finally, we provide an overview of existing benchmark generators for clustering, and analyze their strengths and limitations.

A. The general role of benchmarks

Empirical comparison between techniques is a cornerstone of the scientific method. At a community-level, methods developed by independent researchers need to be compared in order to gain insights into the applicability, generalizability, and efficacy of their developments. As direct comparison is only

possible on the same problem instances, subsequent research is highly likely to adopt instances used by other researchers. The importance of reproducibility further re-enforces the need for a common, accessible set of data that can be utilized to facilitate comparisons across independent studies. This feedback loop results in “standard” benchmarks becoming virtually required to include in experimentation [10]. This requirement is supported explicitly through the creation of benchmark suites — a collection of problems collated and/or created for the purpose of widespread comparison [13], [14].

The issue with this feedback loop is that the community as a whole risks tuning both hyperparameters and algorithmic development to these specific problems [10], [15]. If these popular problems represent a broad-spectrum of real-world challenges, then this is not a negative; analyzing whether these problems adequately cover the space of encounterable problems, however, is difficult if even possible to do in its entirety [13].

To combat this challenge, Hooker [10] argues for “controlled experimentation” i.e. comparing algorithmic performance specifically on a problem characteristic that the research in question is addressing, compared to the “competitive testing” that is encouraged when the same subset of datasets are re-used time and again. This argument is consistent with the implications of the No-Free-Lunch (NFL) theorem for learning, which supports the intuition that no single algorithm is expected to be superior across all problems [16], and that binary statements about algorithm superiority (i.e. “algorithm A is better than algorithm B ”) are only possible within particular problem classes. Controlled experimentation is inherently simpler with synthetic benchmarks, as we have control over the generating mechanism and thus (to varying extents) the properties of the instances. The use of real-world instances for this purpose is possible only if there is an appropriate measure of the problem characteristic and a controlled way to vary it.

B. The algorithm selection problem

In the above, we have introduced the notion of problem characteristics, and the need for these to be varied in benchmark data. This is to ensure appropriate coverage across the space of possible instances, and the ability to appropriately differentiate between the challenges these instances may pose for different algorithms.

Rice [17] formalized these characteristics as “problem features” in the context of the algorithm selection problem (ASP), where the goal is to predict which algorithm from a portfolio is best-suited to a given problem instance based on its problem features. This is premised on the existence of an identifiable relationship between the problem features and problem difficulty for a given algorithm, typically requiring these features to be specific to the problem class. A series of papers by Smith-Miles further extended this framework [12], [18], [19], using an instance space to visualize the interaction between problem features and algorithmic performance [12] and applying this approach to combinatorial optimization [20], [21] and supervised learning [22].

The problem of selecting an appropriate algorithm for a given task is closely related to the “algorithm configuration problem” (i.e. hyperparameter optimization), which is important in both the metaheuristic and machine learning communities [23]–[25]. In this context, the identification of problem features has featured prominently in the form of exploratory landscape analysis, where quantification of different aspects of the fitness landscape informs not only algorithm configuration, but a more general understanding of the suitability of algorithm components to different problem characteristics [14], [26]–[28].

C. Relevance to unsupervised learning

Unsupervised learning aims to identify the natural structure of a dataset in the absence of information about the ground truth. This pattern recognition task is subjective, even for humans [29], [30], explaining the existence of a diverse range of clustering algorithms that make a broad variety of mathematical assumptions about desirable cluster properties [6], [8].

The dilemma stemming from the NFL theorem is thus exacerbated in cluster analysis, with differing clustering algorithms useful for different clustering problems [6], [31], due to unavoidable differences in the underlying formulation [32]. As the inductive bias of each clustering algorithm differs, this fundamentally governs its capabilities: as an example, the popular K-Means assumes hyper-spherical, compact clusters, and this strictly limits its ability to deal with data that violates this assumption.

In consequence, having a representative and diverse collection of benchmark datasets is particularly crucial in cluster analysis, in order to fairly test each algorithm’s strengths and weaknesses, and to avoid any biases towards those methods whose inductive bias may be consistent or correlated with the assumptions of a limited benchmark set. However, compared to other optimization problems, evaluation of this is complicated by the multifarious nature of cluster quality measures itself [6], [8]. Without explicit effort to first identify and quantify aspects of problem structure or difficulty that matter in cluster analysis, it is challenging to assess and ensure appropriate diversity of problem instances and, thus, a comprehensive benchmark suite. Any efforts to improve existing benchmarks must therefore involve such identification and quantification as an integral step.

Previous works on directly tackling the ASP for clustering (in the context of “meta-learning”) have used generic statistical properties of the data (such as the mean of the moments across the data). An alternative is the use of measures directly characterizing a given group structure [33], [34]. The latter approach is likely to be more powerful in identifying drivers of algorithm performance, but its limitation lies in the difficulty of transferring insights on feature-performance correlations to data with an unknown group structure.

D. Existing clustering benchmarks and generators

In the following, we consider existing clustering benchmarks and the extent to which these meet the criteria of representativeness and diversity, as discussed above.

The datasets from the UCI Machine Learning Repository [35] comprise a diverse mix of real-world datasets and are one of the most common benchmarks for the evaluation of machine learning algorithms. Recent analysis demonstrates, however, that even with real-world data from a range of application areas, diversity in complexity is not guaranteed. Macia et al. [13] analyzed the datasets of the UCI Machine Learning Repository, finding a surprising similarity in complexity across them. They discovered a lack of diversity in both the statistical properties of the datasets and the relative classification performance of different algorithms and parameter settings.

At the other end of the spectrum, two-dimensional toy datasets are a common approach to the evaluation of clustering methods [36], [37]. Typically, these datasets have been handcrafted to illustrate simple capabilities or properties of clustering algorithms. They remain popular due to their simplicity and because they enable intuitive visualization of results, rendering them highly effective at clearly exhibiting a particular property and consequent algorithm behaviour (see e.g. the “moons” dataset¹). Despite these advantages, the scenarios that toy datasets depict are often too contrived.

Although synthetic data is commonly generated *ad hoc* for individual work, some generators have been explicitly designed to have broader utility, with more complex properties than toy data. The generator proposed by Qiu and Joe [38] (abbreviated to *QJ*) uses a geometric framework for cluster placement. The measure of separation used (proposed in [39]), provides a measure of the spatial separation between clusters. By adjusting this minimum amount of separation allowed, the user can generate datasets that have different amounts of separation between clusters, providing some loose control over the resulting difficulty of the datasets. To do this, the covariance matrices are iteratively scaled until this minimum separation is achieved. Although useful and geometrically interpretable, this provides a single perspective of cluster structure. Their generator can, however, embed additional complexity through the addition of noise [38]. Despite having several useful parameters to customize the difficulty of the generated problems, the generator (available as an R package²) is not easily extended to incorporate other measures of cluster structure.

Handl and Knowles [11] (abbreviated to *HK*) created two generators (named “gaussian” and “ellipsoidal”) that have been used extensively for the creation of synthetic clusters. At their core, these generators aim to generate clusters that are as compact as possible with either no (“gaussian”) or minimal (“ellipsoidal”) overlap. The “gaussian” generator uses a trial-and-error scheme where Gaussian clusters are randomly generated, and rejected if they overlap with existing clusters. The “ellipsoidal” generator was designed specifically for generating elongated clusters in higher dimensions, and uses an EA to optimize cluster location such that overall variance in the dataset is reduced while penalizing overlap. Despite the popularity of these generators, their design is rigid with many hand-tuned parameters, and without consideration or

¹Example: <https://rdrr.io/cran/clusterSim/man/shapes.two.moon.html>

²<https://cran.r-project.org/web/packages/clusterGeneration/index.html>

easy scope for extension to consider additional or alternative aspects of cluster structure.

In this paper, we therefore propose a more general evolutionary framework for generating new synthetic benchmarks for cluster analysis. Our work tackles two crucial steps for the construction of diverse datasets: (i) the explicit identification of problem characteristics that are relevant to differentiating clustering algorithm performance and (ii) the design of an optimizer that is sufficiently general to evolve benchmarks for a range of criteria. Specifically, we experiment with two different types of criteria: approximating a desired target value of problem characteristics, and directly maximizing performance differences between algorithms. The latter approach helps highlight the strengths and weaknesses of the contestant methods, providing direct insight into their individual inductive bias. To evaluate the performance of our approach, we compare datasets obtained from our framework against multiple generators/dataset collections in terms of both the spread of performance across clustering algorithms, and the variance across our proposed set of problem features. This work greatly extends an initial proof-of-concept outlined in [40], generalizing the framework and introducing the ability to generate datasets directly for particular clustering algorithms.

III. HAWKS

We now proceed to introduce the details of our framework, HAWKS, for the generation of diverse, complex datasets. HAWKS uses an EA to evolve a population of datasets, where the objective function and constraints work together to vary properties of the resulting datasets that affect clustering algorithm performance. For each component of HAWKS, we discuss its general role in dataset generation, and motivate the specific design choice made for the experiments presented in this paper.

A. Representing a dataset

In HAWKS, a dataset is represented by a set of clusters, which are themselves defined by the parameters of a given multivariate distribution. With this level of abstraction, we avoid handling of individual data points, and provide an intuitive manipulation of clusters through adjustment of distribution parameters. In principle, this approach allows for the inclusion of a variety of cluster generating models, given suitable initialization and variation operators for each distribution have been defined (see Sections III-B and III-E).

For the experiments reported in this paper, each cluster is represented by a multivariate Gaussian distribution, due to their relative simplicity and prominence in a machine learning context. For a dataset in D dimensions, each cluster is therefore encoded by a $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ pair, which we will refer to as a single gene. Here, $\boldsymbol{\mu}$ is the D -dimensional mean vector and $\boldsymbol{\Sigma}$ is the symmetric $D \times D$ covariance matrix. A dataset with K clusters is therefore represented as a genotype composed of K genes. In the current implementation, K is set *a priori* as an input parameter.

B. Initializing a population of datasets

The first step of HAWKS is to create an initial population of datasets. The sizes of the clusters can be controlled in two ways: (i) all clusters can be of equal size or (ii) each cluster is of random size (with an optional minimum size to avoid clusters small enough to be considered outliers)³. In both cases, the user predefines the total number of data points N . The cluster sizes remain fixed across all individuals for the remainder of the evolution, to avoid interference with the fitness (described in Section III-C) and focus of the search on the distribution parameters only.

All other aspects of the cluster-level initialization are specific to the type of distribution used; here, we describe our approach for Gaussian clusters. The initial means are sampled from a D -dimensional uniform distribution, i.e. for the i th cluster $\boldsymbol{\mu}^{(i)} \sim U(0, \beta_1)^D$ where β_1 is the upper threshold to control the initial sampling space for the means. In HAWKS, the covariance matrix is defined for the i th cluster as $\boldsymbol{\Sigma}^{(i)} = \mathbf{R}^{(i)} \mathbf{S}^{(i)} \tilde{\boldsymbol{\Sigma}}^{(i)} \mathbf{R}^{(i)\top}$, where $\mathbf{R}^{(k)}$ and $\mathbf{S}^{(k)}$ are the k th rotation and scaling matrices, respectively, and $\tilde{\boldsymbol{\Sigma}}^{(k)}$ is the k th axis-aligned covariance matrix (i.e. a diagonal matrix that consists of only the variances). These variances are sampled similarly to the means, i.e. $\tilde{\boldsymbol{\Sigma}} = \text{diag}(U(0, \beta_2)^D)$ where β_2 is the upper threshold to control the initial sampling space for the variances. The scaling matrix is set to the identity matrix at this stage. The covariances are then obtained via rotation using a random rotation matrix (\mathbf{R}), which is drawn from the Haar distribution [41] to generate a valid covariance matrix. This method permits generation of clusters with a variety of shapes and orientations, thus ensuring the initial population has a diverse set of individuals.

This approach, while more complex than constructing a covariance matrix with random values, allows both more intuitive parameterization and the ability to separately modify individual components of the covariance matrix, which we later exploit for mutation (Section III-E).

C. Computing the fitness of a dataset

The quantification of fitness plays a key role in focusing the search on those datasets that are deemed to present interesting clustering benchmarks. It is thus one of the most vital design decisions in a given generator, but is complicated by our limited understanding of the desirable properties of clustering benchmarks, as discussed in Section II-C. Therefore, our framework is designed to allow for the interchangeable use of different fitness functions, providing scope for a variety of choices to be trialed.

In this paper, we introduce and experiment with two different *modes* of optimization which support the generation of datasets for distinct goals. The first, named *Index* mode, optimizes the datasets towards a user-defined target value for a given cluster validity index. This provides control of the broad difficulty of the datasets, where the nature of the difficulty is governed by the properties of the index used. The second,

³Owing to its non-triviality, the method to generate random cluster sizes with a minimum size that overall sum to N is described in the supplementary material (Section ??).

named *Versus* mode, directly optimizes datasets such that they maximize the performance difference between two clustering algorithms. By generating a range of datasets that are simple and difficult for specific algorithms we can “stress-test” them, revealing their relative strengths and weaknesses. We describe these approaches in more detail in the following sections.

1) *Index Mode*: Cluster validity indices capture the amount of structure of a partition, in terms of the underlying data distribution. Given knowledge of the generating model (i.e. the true partition), they can thus act as a proxy for the ease of recognizing this structure in a given dataset.

As previously discussed, there are many different cluster validity indices, each with a slightly different perspective of how cluster structure should be quantified. Previous work, including [9], could not conclude superiority of any single validity index. Thus, to obtain a comprehensive understanding of the difficulty of a given dataset, many cluster validity indices could potentially be used in combination [42]. In the context of a fitness function, this could be done in the form of an aggregation of indices or through formulation as a many-objective problem. Here, we opt for a middle ground, limiting ourselves to the definition of fitness through a single cluster validity index, but allowing for the incorporation of additional considerations (which could include validity indices) through the use of constraints (discussed in Section III-D).

For the experiments reported in this paper, we use the silhouette width [43] as a representative example of a validity index. It is an established, widely-used method and has two particular characteristics that makes it a promising choice for use here: (i) it is a very rich measure that provides information at multiple level of resolutions (the individual data point, the cluster, and the partition-level); and (ii) as it is bounded in the range $[-1, 1]$, it can be easily compared across individuals and even runs (when a similar dimensionality is used), and it is readily interpretable to the user.

The silhouette width is a combination internal validity index, as it measures a ratio of intra-cluster compactness and inter-cluster separation [1]. For a single data point, \mathbf{x}_i , it is defined as:

$$s(\mathbf{x}_i) = \frac{b(\mathbf{x}_i) - a(\mathbf{x}_i)}{\max\{a(\mathbf{x}_i), b(\mathbf{x}_i)\}}. \quad (1)$$

Here, $a(\mathbf{x}_i)$ represents the cluster compactness (with respect to \mathbf{x}_i) and is the average distance from \mathbf{x}_i to all other data points in its cluster. The separation between clusters is represented by $b(\mathbf{x}_i)$; for data point \mathbf{x}_i this is defined as the minimum of the average distances to all data points in every other cluster. The silhouette width is calculated for all N data points in dataset X , and an average is taken to obtain the overall silhouette width:

$$s_{all} = \frac{1}{N} \sum_{i=1}^N s(\mathbf{x}_i). \quad (2)$$

A value of 1 represents very compact and well-separated clusters, whereas a negative silhouette width value indicates that points in different clusters are not well-separated (and that their cluster membership should be changed).

Independent of the validity index chosen, direct maximization or minimization of such an index would always lead to

the evolution of datasets that are trivially separable or fully unstructured, respectively. HAWKS therefore requires input in the form of a desired target value, allowing direct modulation of the desired level of structure.

In other words, and using the example of the silhouette width, a target value (denoted s_t) is specified by the user and datasets are then optimized to meet this target value. This is achieved by minimizing the absolute difference between s_t and s_{all} , defined as:

$$\min f(\boldsymbol{\mu}^{(1)}, \boldsymbol{\Sigma}^{(1)}, \dots, \boldsymbol{\mu}^{(K)}, \boldsymbol{\Sigma}^{(K)}) \equiv \min |s_t - s_{all}|. \quad (3)$$

2) *Versus Mode*: The *Index* mode provides us with the ability to generate benchmarks that meet specific thresholds of user-defined validity criteria. However, shared assumptions between validity indices and clustering algorithms themselves can make it difficult to generate datasets with properties that specifically challenge a given algorithm. This is particularly the case when working with clustering techniques whose inductive biases are poorly understood, e.g. self-organizing approaches or those that utilize deep learning [44].

To perform “controlled experimentation”, a more direct link between dataset generation and algorithm performance may be needed [10]. Our *Versus* mode tries to address this by directly optimizing the performance difference between two algorithms, thereby exploiting the strengths of one algorithm, the weaknesses of the other, or a combination of the two.

We re-formulate the fitness function such that we maximize the difference between the scores of the ‘winning’ algorithm (\mathcal{A}_w) and the ‘losing’ algorithm (\mathcal{A}_l), using a scoring function ϕ :

$$\max f(\boldsymbol{\mu}^{(1)}, \boldsymbol{\Sigma}^{(1)}, \dots, \boldsymbol{\mu}^{(K)}, \boldsymbol{\Sigma}^{(K)}) \equiv \max(\phi(\mathcal{A}_w) - \phi(\mathcal{A}_l)). \quad (4)$$

While any scoring function can be used, we have access to the generating model of the clusters and therefore the ground truth. The quality of each partition can therefore be objectively assessed using an external validity index such as the Adjusted Rand Index (ARI) [45]. The ARI measures the co-occurrences of cluster assignments between two partitions, which in our case corresponds to the output of a given clustering algorithm, and the ground truth. The upper bound of 1 represents identical assignment, whereas 0 indicates random assignment. As the ARI considers only the assignment of points, it can be used with any clustering algorithm and has no preference of structure.

D. Augmenting cluster properties using constraints

As previously discussed, it is difficult to define a single fitness measure that represents all properties that may be desirable in a benchmark dataset. Therefore, our framework uses constraints to allow for the integration of additional considerations, with two main aims: (i) to avoid the generation of trivial datasets that are e.g. simple for any algorithm, or too noisy to be clustered, and (ii) to introduce additional properties that balance limitations of individual fitness functions or further enhance diversity in the datasets obtained. To control optimization of the fitness and constraints, we use stochastic

ranking [46] to balance the satisfaction of all constraints (this is discussed further in Section III-F).

In the following, we discuss two constraints introduced to meet these aims by directly accounting for local overlap between clusters and for cluster shape (as measured by cluster elongation — for a different generating model, other choices may be plausible). These choices of constraints have to be seen in the context of the fitness functions adopted for our experiments: when using the silhouette width in *Index* mode, the fitness provides a powerful overall perspective of separation between clusters. However, the reliance on averaging can lose fine-grained information about local variability, e.g. lack of inter-cluster separation for small clusters. Furthermore, the silhouette width does not directly capture information about cluster shape, while differences in cluster shapes are known to drive some performance differences between algorithms. Similarly, for the *Versus* mode, a performance difference between the algorithms is sought with no explicit concern for the underlying cluster properties. While this exploratory approach is desired, there are situations where this can lead to cluster structures that are not meaningful (e.g. clusters that are completely overlapping) but one algorithm is still perceived to be favoured due to arbitrary artefacts.

1) *Overlap*: Real-world data typically does not contain cleanly separated clusters, either due to noise or reflective of the underlying relationships between variables. Clustering algorithms differ significantly in their ability to handle some degree of cluster overlap, and control over this aspect of our benchmark is therefore important. At the extreme ends, a benchmark set with very well-separated clusters may be too simple to detect any performance differences between algorithms; equally, a dataset where clusters fully overlap will be of no use.

The definition of overlap is not purely objective, and can be implemented in different ways. Fränti and Sieranoja [36] defined a data point as overlapping if it is closer to a centroid of a different cluster than to the centroid of its own cluster. This definition, however, does not extend to highly eccentric clusters. We use a definition similar to [11], where a data point is considered as overlapping if its nearest neighbour belongs to a different cluster. In addition to avoiding the compactness bias introduced by using centroids, this definition has the specific benefit of countering an inherent limitation of the silhouette width, where a high average silhouette width value can be driven by having large clusters that are very well-separated and fails to sufficiently reflect the presence of very small but highly overlapping clusters.

We calculate the overlap as the percentage of data points whose nearest neighbour is in a different cluster:

$$overlap = 1 - \frac{1}{N} \sum_{\mathbf{x} \in X} \mathbb{1}_{C_k}(n_{\mathbf{x}}) \quad (5)$$

where C_k is the cluster that data point \mathbf{x} belongs to, $n_{\mathbf{x}}$ is the nearest neighbour of \mathbf{x} , and $\mathbb{1}(\cdot)$ is the indicator function that is 1 if $n_{\mathbf{x}} \in C_k$ and 0 if $n_{\mathbf{x}} \notin C_k$.

2) *Elongation*: The elongation (or, more specifically, the eccentricity) of clusters can pose problems for compactness-based algorithms such as K-Means, and thus presents an

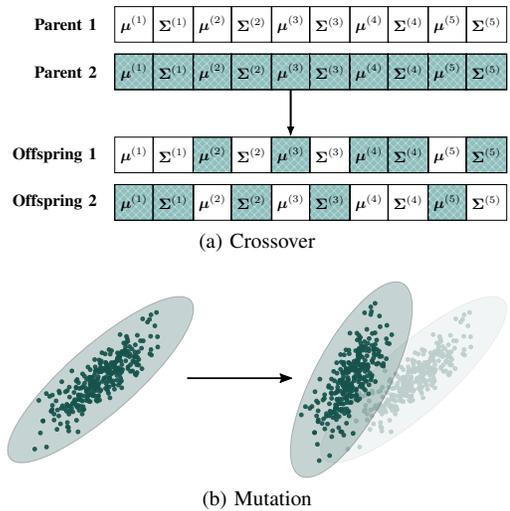


Fig. 1. Illustrations of the HAWKS genetic operators. Uniform crossover between two individuals is shown in (a), where the means and covariances are swapped independently. In (b), a single cluster is mutated where both the location (mean) and shape (covariance) have been randomly perturbed (the original cluster is shown again faded to illustrate the relative differences).

additional aspect that we are keen to control. Although our initialization parameters can adjust initial cluster eccentricity, by using an explicit constraint we can encourage (or penalize) this further during the evolution.

Our definition of this constraint is specific to the generating model used, and different definitions are possible. As previously mentioned, each full covariance matrix (Σ) in our cluster representation is separated into the axis-aligned variances ($\tilde{\Sigma}$) and a rotation matrix. As the variances on the diagonal of $\tilde{\Sigma}$ are the eigenvalues of the full covariance matrix, i.e. $\tilde{\Sigma} = \text{diag}(\lambda_1, \dots, \lambda_D)$, the ratio of the maximum and minimum of these eigenvalues gives us a measure of the cluster eccentricity. As even a single eccentric cluster can pose challenges for compactness-based algorithms, we take the minimum of these ratios across all K clusters, i.e.:

$$\lambda^{ratio} = \max_{\forall k \in \{1, \dots, K\}} \frac{|\lambda_{\max}(\Sigma^{(k)})|}{|\lambda_{\min}(\Sigma^{(k)})|}, \quad (6)$$

where $\lambda_{\max}(\Sigma^{(k)})$ and $\lambda_{\min}(\Sigma^{(k)})$ are the maximum and minimum eigenvalues of $\Sigma^{(k)}$, respectively.

E. Perturbing a dataset

As is typical in an EA, the variation operators in HAWKS provide a further opportunity to integrate domain knowledge and focus the search. Specifically, a close alignment between the generating model used (i.e. the representation of an individual cluster and partition) and the operators is crucial to optimization performance and the types of datasets that can be obtained. In the following, we describe the variation operators designed for the multivariate Gaussian distributions utilized as the generating model in our experiments.

1) *Crossover*: For recombination between datasets, HAWKS uses a high-level uniform crossover scheme where the two components defining each cluster distribution (i.e. μ and Σ) can be swapped separately between individuals. This

allows for both the location (mean) and shape (covariance) of clusters to be swapped between individuals independently, as illustrated in Fig. 1a.

2) *Mutation*: Meaningful mutations of a cluster require careful design of a geometrically relevant operator that is relevant for the distribution being used. In our Gaussian case, we use a separate operator for the mean and covariance terms.

A key issue identified in [40] was the increasing amount of drift of the mean operator in higher dimensions. The original operator shifts the mean to a random nearby point; at higher dimensionality there are an increasing number of directions that point away from other clusters, and a random walk is thereby likely to increase the silhouette width at most steps. In the supplementary material (Section ??), we test multiple new mutation operators (drawing inspiration from operators used in particle swarm optimization and differential evolution) to address this issue by directly considering the location of other clusters.

The resulting operator selected and used throughout this paper utilizes concepts from particle swarm optimization [47], where a combination of other centroids and a global representative is used to embed direction either towards or away from existing clusters, directly affecting the fitness. The new mean, $\mu^{(i)}$, is obtained as follows:

$$\mu^{(i)} = \begin{cases} \mu^{(i)} + [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } p \leq 0.5 \\ \mu^{(i)} - [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } p > 0.5, \end{cases} \quad (7)$$

where $\mu^{(i)}$ is the current cluster mean, $\mu^{(n)}$ is the mean of another randomly selected cluster, $\bar{\mu}$ is the global mean across all data points, w_1 and w_2 are random weighting coefficients in the range $[0, 1]$, and p is a random coin-flip to decide whether to move away from or towards this weighted combination of an existing cluster mean and the global mean.

The covariance governs the shape of the cluster, affecting both the overlap between clusters and the capabilities of compactness-based algorithms (such as K-Means). To mutate the covariance, we rotate the cluster and scale its eigenvalues, effectively changing the shape of the cluster. Similar to the initialization, the rotation matrix is drawn from the Haar distribution [41], though here it is raised to a fractional power to avoid complete reorientation of the cluster. The scaling matrix, S , is drawn from a Dirichlet distribution in order to ensure that the resulting determinant is unchanged i.e. $\det(\tilde{\Sigma}) = \det(S \cdot \tilde{\Sigma})^4$. Using an analogy, this has the same effect as rotating a balloon and applying pressure to the principal semi-axes, thereby changing the shape while maintaining the volume. The combined effect of both mutation operators is illustrated in Fig. 1b.

F. Selecting a dataset

As previously discussed, we use stochastic ranking [46] to balance the satisfaction of the objective and constraint(s), which may be complementary or impossible to both fully satisfy. By adjusting the probability of comparing two individuals on their fitness (denoted P_f), we can effectively

weigh the satisfaction of the objective or the constraints, providing us with another way of controlling the properties of the resulting datasets. For example, if using only the overlap constraint in *Index* mode, setting P_f to a higher value will add selection pressure to datasets with a silhouette width closer to s_t , potentially at the cost of a higher degree of overlap. Thus, unlike the traditional use of stochastic ranking that uses a narrow range of values for P_f [46], [48] to avoid too much weighting towards the infeasible or feasible regions, we can use the entire range (as datasets that heavily violate the constraints may still be useful).

Similar to the original work [46], for environmental selection we use stochastic ranking to select the top $|\mathcal{P}|$ individuals from the sorted pool of parents and offspring. For parental selection, we use standard binary tournament where the rank is used to determine the winner (in lieu of using only the fitness) to ensure a continued selection pressure towards individuals that best satisfy the fitness and constraints as weighted by P_f .

IV. EXPERIMENTAL SETUP

In order to assess the capabilities of HAWKS, we have three primary aims in our experiments: (i) compare the diversity of performance of multiple, distinct clustering algorithms on datasets from HAWKS against that seen for other popular generators or dataset collections; (ii) compare the diversity across a distinct set of problem features, to see if datasets from HAWKS cover a wider space of properties than other datasets; and, (iii) gain insights into the algorithms themselves, utilizing the *Versus* mode of HAWKS to directly challenge clustering algorithms.

The remainder of this section describes our proposed problem features (Section IV-A), the generators and dataset collections we compare against (Section IV-B), and the relevant parameters HAWKS used in each experiment (Section IV-C).

A. Problem features

In order to measure dataset diversity (with respect to their inherent properties), we need to define a set of problem features describing properties relevant to algorithm performance. This is central to the algorithm selection problem (ASP), as these features are used to predict which algorithm is best for a given problem instance. These problem features are also vital for the creation of an instance space that visualizes the datasets, allowing identification of areas in the space where particular algorithms are most suitable.

In previously discussed work on the ASP for clustering, the problem features (also referred to as “meta-features”) used are generally statistical or information-theoretic, and not specific to clustering [33], [34]. The problem features previously used in [40] were also too simplistic, and fully coincided with parameters available in HAWKS, reducing the utility of the instance space.

As there are many properties that influence performance for a given clustering algorithm, and for most algorithms this aspect is not fully understood, a full complementary set of problem features is arguably impossible [6], [9]. Nonetheless, going beyond simple statistical measures of the data by

⁴This is described in further detail in Section ??.

incorporating measures specific to clustering (especially those that make use of the generating model) should improve the discriminative power of the problem features.⁵ The remainder of this section describes our proposed set of problem features for this task.

1) *Average eccentricity*: A high amount of cluster eccentricity can cause problems for compactness-based algorithms. To measure this, we use a similar method to its calculation as a constraint in HAWKS, with some key differences. For both the synthetic and real-world datasets used in this paper, we have the true labels. Using these, the eccentricity is calculated as before — calculating the ratio of the maximum to minimum eigenvalues for each cluster, except that the average (rather than minimum) value is used as the problem feature. The eigenvalues are obtained via singular value decomposition, but (in contrast to the constraint) a subset of the eigenvalues is used. This reduces the sensitivity to outliers and subspace clusters (which result in zero, or arbitrarily close to zero, eigenvalues). For further details, see Section ??.

2) *Connectivity*: The connectivity, defined by Handl and Knowles [49] (which is a modified version of what was proposed in [50]), measures the extent to which neighbouring data points are assigned to the same cluster. We normalize this measure by the number of data points to get a measure that is comparable across datasets. This provides a more nuanced picture than the overlap constraint used by HAWKS, which looks at the single nearest neighbour rather than the set of L nearest neighbours, where L is a parameter of the connectivity. Here, we use $L = 10$ in line with previous work [49].

3) *Dimensionality*: As different measures of distance or similarity are affected differently by high-dimensionality, this feature simply describes the number of dimensions in the data. This is our only feature that is ground-truth agnostic, and is shared with previous work on the ASP for clustering [33], [34].

4) *Entropy of cluster sizes*: The density or relative sizes of the clusters can affect the performance of different algorithms. For example, the ability of K-Means to discover clusters is diminished when a small subset of clusters contain the majority of data points. We calculate the entropy of cluster sizes as follows:

$$H(\mathcal{C}) = - \sum_{k=1}^K \frac{|C_k|}{N} \log_K \left(\frac{|C_k|}{N} \right), \quad (8)$$

where \mathcal{C} is the set of K clusters, and $|C_k|$ is the cardinality of cluster C_k , which is normalized by the size of the dataset (N). We use \log_K to compare across datasets such that, for any K , a perfectly equal distribution of cluster sizes results in $H(\mathcal{C}) = 1.0$, and $H(\mathcal{C}) \rightarrow 0$ when one cluster has $N - K + 1$ data points.

5) *Number of clusters*: The number of clusters can inherently affect algorithms where initialization of cluster location is key to convergence, particularly if clusters are well-separated.

6) *Silhouette width (average)*: The average silhouette width measures how similar the data points are within their own clusters, and how well-separated these are to the other clusters, indicating the potential difficulty for clustering algorithms in discovering these clusters.

7) *Silhouette width (standard deviation)*: Averaging the silhouette width can obscure the presence of a small number of very well-separated clusters, and ill-defined overlapping clusters. Thus, the standard deviation of the silhouette width (calculated across all data points) indicates whether all points are similarly separated. Higher values indicate the presence of overlapping clusters (from a different perspective than the connectivity measure), which clustering algorithms have different capabilities in handling.

B. Other generators and datasets

To provide context for the diversity in performance and problem features for the datasets produced by HAWKS, we compare against multiple generators or collections of popular datasets that have been previously used to evaluate clustering algorithms. We briefly describe these below, though more details about the parameters used for these datasets, and how they compare (in terms of their size, dimensionality etc.) can be found in the supplementary material (Table S-??) and their respective papers.

1) *HK*: This is a collection of 350 datasets generated using the “ellipsoidal” generator proposed in [11], and used as a benchmark in [51]. For 35 unique combinations of the number of clusters and dimensionality, 10 datasets are generated.

2) *QJ*: We use the set of 243 datasets generated using the parameters proposed by Qiu and Joe [38]. The authors calculated three target separation values using their measure of separation: “close structure”, “separated”, and “well-separated”. Further complexity to these datasets is added by specifying varying proportions of noisy variables, alongside small variation in the number of dimensions.

3) *SIPU*: Fränti and Sieranoja [36] introduced a benchmark consisting of multiple sets of clustering datasets, of which we use the ‘S-sets’, ‘A-sets’, and ‘G2 sets’. These sets were originally intended to stress-test compactness-based algorithms, of which K-Means exhibited a wide range of performance. The ‘S-sets’ are all 2D data where $N = 5000$ and $K = 15$, but have different degrees of overlap between the clusters (determined by the aforementioned method of the closest centroid). The ‘A-sets’ are also 2D data with varying numbers of (equally-sized) clusters. The ‘G2 sets’ of datasets consist of two Gaussians with varying degrees of overlap, constant size ($N = 2048$) and varying dimensionality (chosen across the range $2^1, 2^2, \dots, 2^{10}$). The variance of overlap and high number of dimensions in this benchmark presents a variety of challenges for clustering algorithms.

4) *UCI*: The UCI Machine Learning Repository [35] is a popular source of datasets used for machine learning. As noted in [8], [36] the class labels do not necessarily translate to meaningful cluster labels, and thus these datasets may not be the most suitable for clustering. Nonetheless, they have seen extensive use in the clustering literature, and we include them

⁵As discussed earlier, use of the ground truth in these measures comes at the cost of limiting applicability to datasets where such knowledge is unavailable; this is a major issue for ASP applications but is not our intended focus here.

for completeness. Specifically, we use the subset of 20 datasets used by Arbelaitz et al. [9].

5) *UKC*: These 8 real-world datasets, curated in [51], are the (anonymized) locations of different crimes. Alongside the *UCI* datasets, these will help provide insight into whether there are significant differences between the real-world and synthetic data (in terms of their problem features or clustering algorithm performance).

C. HAWKS configurations

In this section we highlight the key parameters for HAWKS, separating those that are common across all experiments and those adjusted for different modes. In particular, while the core EA parameters (population size, mutation probability etc.) remain the same across modes, in our experiment for the *Index* mode we vary the objective function target and constraint parameters to generate different datasets. In contrast, the constraint parameters remain fixed for the *Versus* mode experiment as the variation comes from selecting different pairs of clustering algorithms. The full configurations for both sets of experiments can be found on GitHub⁶.

1) *Common parameters*: In both *Index* and *Versus* mode experiments, the following core EA parameters are used: $G_{\max} = 100$ generations, $\mathcal{P} = 10$ individuals, $p_c = 0.7$ crossover probability, and $p_m = 1/K$ mutation probability for the mean and covariance mutation operators. The low population size was previously found to be sufficient for both diversity and convergence [40].

2) *Solution selection*: At the end of a run, a single dataset needs to be selected. Here, we simply select the individual with the highest fitness. Different choices are possible, e.g., due to our use of stochastic ranking, the individual with the highest (sorted) rank could be selected instead (though this had little effect on the experiments in this paper).

3) *Analyzing benchmark dataset diversity (Index mode)*: To test the ability of HAWKS to produce a variety of datasets, we vary several parameters to encourage coverage of our feature space:

- A poor and a high silhouette width ($s_t \in \{0.45, 0.9\}$).
- Two upper thresholds for the *overlap* constraint ($overlap \leq \{0, 0.1\}$) to penalize any overlap, or penalize if more than 10% of the data points overlap.
- Two lower thresholds for the eccentricity constraint ($\lambda^{ratio} \geq \{1, 50\}$) to allow for any amount of eccentricity, or encourage all clusters have some eccentricity.
- Two levels of dimensionality ($D \in \{2, 50\}$), dataset size ($N \in \{500, 2500\}$), and number of clusters ($K \in \{5, 30\}$).

Of note is the use of $P_f = 0.5$, which weights the fitness and constraints equally. By varying s_t , we directly attempt to generate datasets that either have poor cluster separation or are well-separated. Different values of the overlap and eccentricity constraints help further modulate the level of separation and the minimum cluster eccentricity. HAWKS is run 7 times for each of the 64 unique combinations of parameters listed above, resulting in 448 datasets.

As we aim to measure diversity in clustering performance, we need an objective way to measure this. We therefore use the ARI (see Section III-C2) to compare the ground truth (which we know for all datasets used here) with the assignments from each algorithm. For the clustering algorithms themselves, we select four well-established algorithms with distinct properties and inductive biases, allowing us to assess the diversity of challenges that the datasets pose. These are:

- Average-linkage — A hierarchical clustering method that uses the average distance between all points in a cluster when deciding which are the closest (and thus should be merged).
- Gaussian mixture models (GMM) — Probabilistic models which try to represent subpopulations (of the data) through a number of Gaussian distributions. To obtain a crisp clustering, each point is assigned to the cluster with the highest probability.
- K-Means++ — Proposed by Arthur and Vassilvitskii [52] to improve K-Means, the initialization scheme probabilistically selects cluster centres such that points that are further away from existing cluster centres are more likely to be selected.
- Single-linkage — In contrast to average-linkage, single-linkage considers the distance between two clusters as the shortest distance from *any* member of one cluster to any member of another.

To assess the maximum potential of each algorithm, we provide the true number of clusters for each dataset. Owing to the propensity of the linkage algorithms (particularly single-linkage) to be side-tracked by singleton clusters, we also run average- and single-linkage with double the true number of clusters [53].

4) *Challenging specific algorithms (Versus mode)*: To ascertain the capabilities of HAWKS' *Versus* mode, we run each of the four clustering algorithms against the others in a head-to-head, with either as the 'winner' and 'loser', respectively.

Some pairings of algorithms (e.g. K-Means++ vs. GMM⁷) are likely to be more competitive, due to shared capabilities and inductive biases. We expect this to translate into a reduction of the performance differences that can be observed (and thus a weaker fitness gradient). While the constraints are still important in such a scenario, we wish to avoid situations where the optimization is being mainly driven by reducing the constraint penalty. Rather than simply removing the constraints, we therefore increase P_f to emphasize the production of a difference in algorithmic performance. This provides a useful lever for weighting the maximization of the performance difference against producing datasets that do not sacrifice cluster structure (e.g. through a high overlap). As such, we use $P_f = 0.75$ in all head-to-head runs. Further adjustment of this parameter is encouraged for complex algorithm pairings.

Finally, all experiments presented for the *Versus* mode are run using $K = 5$ clusters and $N = 2000$ data points. The dimensionality is consistently set to $D = 2$, ensuring a straightforward visualization of the datasets and enabling us

⁶https://github.com/sea-shunned/hawks_configs

⁷Referring to Eq. 4, we use the format \mathcal{A}_w vs. \mathcal{A}_l .

to observe the properties of these datasets without concerns about information loss through projection.

V. EXPERIMENTAL RESULTS

In this section, we explore HAWKS’ ability to produce datasets that exhibit a broad range of properties and pose challenges to different clustering algorithms. Separate results for the *Index* and *Versus* modes highlight the general flexibility of our framework.

A. Analyzing benchmark dataset diversity (Index mode)

This section presents the instance space constructed from our set of 7 problem features applied to the 1,176 datasets from 6 popular collections/generators and our generator HAWKS. We assess diversity by considering (i) the variation observed across problem features and (ii) the variation observed in the performance of 4 clustering algorithms.

First, we use the instance space to understand how the datasets are spread across the space, and if there are distinct patterns either in the sources of the datasets or in the performance of the clustering algorithms. The two principal components of the instance space shown in Fig. 2 account for 57.56% of the variance, and so there is some information loss in the projection. Nevertheless, it is evident that the components capture a sufficient proportion of the variance to highlight some key differences between datasets.

The instance space in Fig. 2a uses colour and marker coding to flag up which collection each dataset comes from. Visualizations of how each problem feature varies across the space can be found in Fig. S-???. The distinct spread of HAWKS datasets across the central part of the space is encouraging and highlights an appropriate level of diversity in terms of the problem features. Furthermore, the interpretation of the principal components (in terms of the underlying problem features), provides clear guidance on additional experiments that could be conducted to expand coverage in various directions. In contrast, the *QJ* datasets expand across a narrow band, indicating a lack of variance across the problem features. The *SIPU* datasets show a strong banding of instances, indicating that there is little variance among the datasets from each configuration (the separated datasets at the bottom of the space are the ‘G2 sets’, which have a much higher dimensionality than other datasets). The *UCI* datasets are spread across the upper-half of the space, though this is unfortunately due to a lack of structure (which we later explore when looking at clustering algorithm performance), as their higher connectivity indicates that the labels do not line up with a spatial perspective of clustering (shown in Fig. S-??). The *UKC* datasets do not seem to represent anything extraordinary with regards to the problem features we use here, leading to the conclusion that either the synthetic datasets used here are not too dissimilar to real-world data or our set of problem features does not capture some aspect of complexity that they uniquely exhibit. Notably, the *HK* datasets are distinct from every other collection, indicating that they have unique characteristics. As shown in Fig. S-??, the main difference is the average eccentricity of the clusters, which is higher than in any other collection. As these datasets

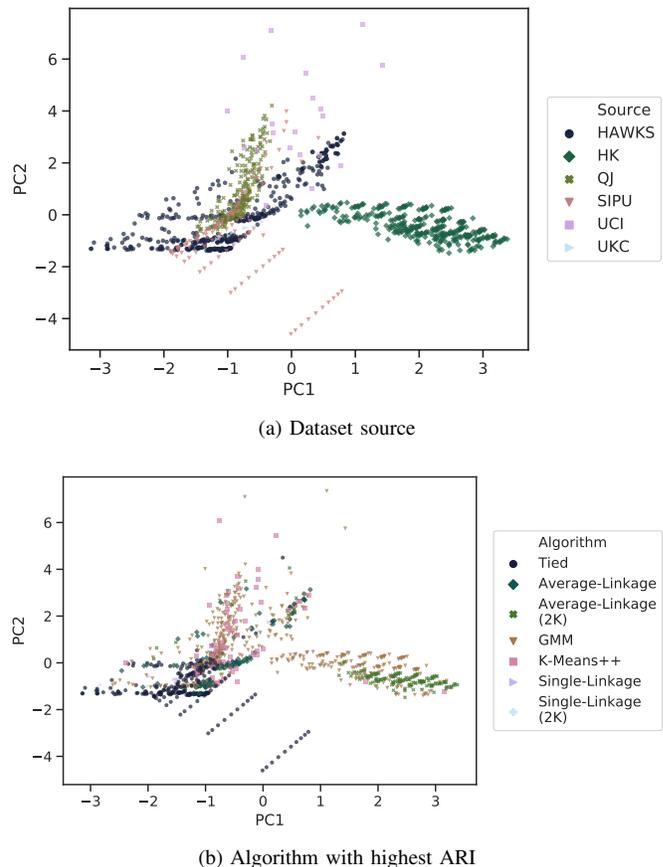


Fig. 2. Two instance spaces, with the source of the dataset (a) or clustering algorithm with the highest ARI (b) highlighted. The 7 problem features are projected down to 2D using PCA. The contribution of each problem feature to the two principal components is $\begin{bmatrix} -0.029 & 0.208 & 0.519 & 0.308 & 0.506 & -0.256 & 0.520 \\ 0.675 & -0.355 & -0.206 & -0.007 & -0.120 & -0.555 & 0.231 \end{bmatrix}$ for the connectivity, dimensionality, average eccentricity, entropy, number of clusters, silhouette width (average), and silhouette width (standard deviation) respectively.

originate from the ‘‘ellipsoidal’’ generator, which was designed exclusively to create eccentric clusters in higher dimensions, this is unsurprising. Of interest, however, is whether this leads to a distinct difference in the relative performance of clustering algorithms on that benchmark suite.

To consider this aspect, Fig. 2b shows the instance space again, but with the clustering algorithm achieving the highest ARI for a given dataset highlighted. The ‘tied’ category is used when at least two algorithms were able to achieve the same ARI, which in every case was 1 and thus the dataset was trivial to cluster. This visualization permits easy identification of potential footprints, though the full information is tabulated in Table I, for completeness. It is evident that there are some areas of the space that correlate with a high performance of a particular algorithm, and this happens for the *HK* benchmark suite in particular, where the highly ellipsoidal clusters consistently favour GMM or average-link with the higher setting of K .

In the more central part of the instance space, however, such footprints are unclear, indicating a complex performance landscape. This could be a consequence of the projection

TABLE I
NUMBER OF TIMES (AS A PERCENTAGE OF THE TOTAL NUMBER OF DATASETS FOR EACH SOURCE) EACH ALGORITHM ACHIEVED THE HIGHEST ARI FOR A GIVEN DATASET. THE BEST PERFORMING ALGORITHM ON EACH DATA SOURCE IS HIGHLIGHTED IN BOLD.

Source	Average-Linkage	Average-Linkage (2K)	GMM	K-Means++	Single-Linkage	Single-Linkage (2K)	Tied
HAWKS	0.261	0.056	0.239	0.056	0.016	0.011	0.362
<i>HK</i> [11]	–	0.426	0.569	0.006	–	–	–
<i>QJ</i> [38]	0.008	0.021	0.646	0.082	–	0.004	0.239
<i>SIPU</i> [36]	–	–	0.106	0.202	–	–	0.692
<i>UCI</i> [35]	0.100	0.150	0.500	0.200	–	–	0.050
<i>UKC</i> [51]	–	–	0.625	0.124	–	–	0.250

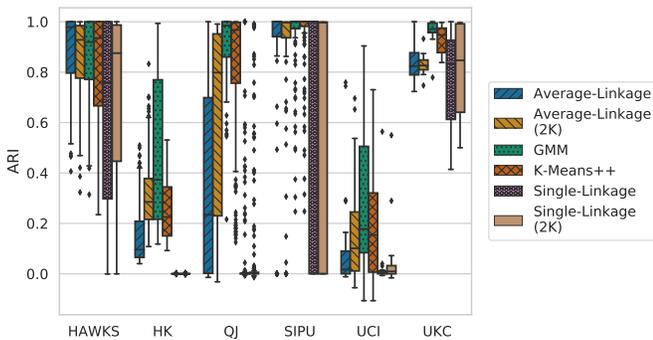


Fig. 3. Clustering performance for each algorithm for each set of datasets.

(which may lose too much information to fully distinguish the datasets), or may point to the role of other features that have not yet been captured here, but distinctly impact on performance. Notably, variation in the identity of the best performing algorithm is the most pronounced for the HAWKS benchmark (see Table I), which is the only collection of datasets where every algorithm was best for a particular dataset. As this was achieved through just a few parameter settings, this is encouraging for the potential of HAWKS to generate diverse datasets.

To further examine the performance variation of the clustering algorithms across these datasets, Fig. 3 shows boxplots for the aggregated performance of each algorithm for each group of datasets. Here, larger boxplots indicate a greater variety of performance for that clustering algorithm, which is preferred. For HAWKS, the boxplots indicate that its datasets elicit a broad range of performance across all algorithms, though the high median ARIs for all algorithms indicate that in general the datasets were not that difficult. As we discourage overlap and half of the datasets were optimized to have a high silhouette width ($s_t = 0.9$), this is somewhat expected and could be addressed by revised parameter choices.

The near-perfect average performance for all algorithms but single-linkage on the *SIPU* datasets indicate their relative simplicity. Similarly, the poor average performance across the *UCI* datasets is consistent with the low connectivity and silhouette width values previously observed in the instance space, and points to weakly defined structure of the ground truth. The *HK* datasets show a reasonable diversity of performance, and the significantly lower mean ARI indicates that these are much

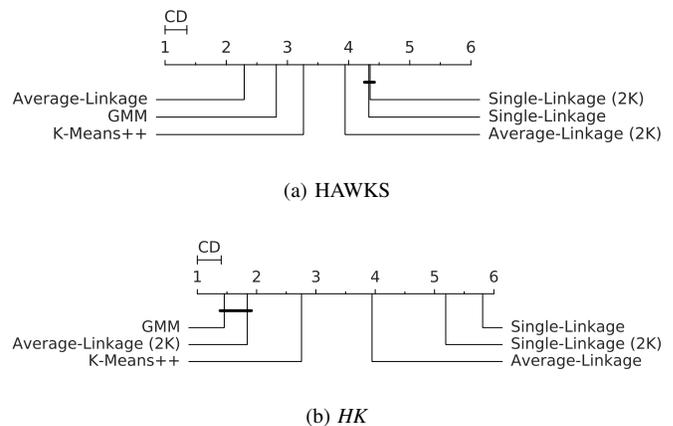


Fig. 4. Critical difference (CD) diagrams, showing the mean rank (in terms of ARI) for the HAWKS and *HK* datasets, for each algorithm. Algorithms connected by solid lines are not significantly different according to a two-tailed Nemenyi test. CD diagrams for the other dataset sources can be found in Fig. S-??.

harder datasets. This may be in part due to the much higher eccentricity of these clusters, but also potentially due to a greater variance in the silhouette width (Fig. S-??), indicating that many data points on the edges of clusters are closer to the points in other clusters than to points in their own. The high performance of all algorithms on the *UKC* datasets indicates that the clusters in this real-world dataset are generally well-defined. This is consistent with the instance space which provided no evidence of unusual complexity.

To include considerations of significance into our analysis, we follow the approach outlined by Demšar [54] to compare multiple methods across multiple datasets. In brief, a Friedman test is used to rank each competing algorithm for each dataset, where the null hypothesis is that all algorithms have equal ranks. As rejection indicates at least one algorithm is significantly different, the (two-tailed) Nemenyi test [55] is used as a post-hoc test to ascertain *which* algorithm is different by calculating the critical difference (CD), which is the minimum that two average ranks must differ by to be significantly ($p < 0.05$) different. We illustrate the results using CD diagrams, which show the average rank of each algorithm across all datasets, with solid lines connecting algorithms whose difference in rank is less than the CD. Well-ordered rankings of algorithms indicate a lack of variance in performance (as an algorithm is consistently bad or good), whereas if all algorithms had an average rank of 3.5 (and thus clustered in the middle of the CD diagram), this would show that the datasets have an equal spread of difficulty for these clustering algorithms.

We show the CD diagrams for the HAWKS and *HK* generators in Fig. 4 (as these two generators showed the greatest diversity in the boxplots)⁸. As evident from the CD diagram, the ranks of the algorithms are more similar for HAWKS, whereas there is a clearer superiority for a subset of algorithms with the *HK* datasets. The best-performing algorithm for HAWKS was average-linkage, which highlights the variety of

⁸The remaining CD diagrams can be found in Section ?? of the supplementary material.

cluster structures that can be generated (as the representation uses purely Gaussian clusters, one might have expected that GMM would on average perform the best). The eccentricity of the generated clusters is reflected in the higher average rank of GMM over K-Means++. Furthermore, the poor performance of single-linkage (using both K and $2K$) indicates that there is (on average) insufficient cluster separation to avoid the ‘chaining’ effect [53], where there is at least one pair of points in different clusters that are closer than a pair of points within a cluster, thereby forming a ‘bridge’ between clusters.

As shown in Table I, all of the generators struggle to produce datasets that single-linkage is uniquely better at, suggesting a possible lower utility of this algorithm, in general. To further investigate this, and to create a fully comprehensive benchmark, it would be of interest to more directly generate datasets that favour single-linkage. In contrast to existing generators, the *Versus* mode in HAWKS facilitates this direct generation, and the results from this mode are discussed in the next section.

B. Generating datasets that challenge specific algorithms (*Versus* mode)

In this section, we present the results of running HAWKS in *Versus* mode, such that we evolve datasets to directly maximize the performance difference between pairs of algorithms. As seen in Section V-A, most current generators have some bias towards a particular algorithm, and no generator (except HAWKS, though not consistently so) was able to produce datasets that were uniquely suited to single-linkage.

First, we need to look at the broad capability of the *Versus* mode in producing a performance differential in the various head-to-heads. In Fig. 5a, we can see a grid of plots showing the performance (ARI) of every algorithm against every other. Each line (in the off-diagonal plots) represents the best dataset from a single run, showing the ARI for ‘winning’ (left) and ‘losing’ (right) algorithms. Here, the angle of the lines indicate the magnitude of the performance difference, and the spread shows the consistency of HAWKS across runs. The plots on the diagonal aggregate the performance for that particular algorithm, e.g. the bottom-right plot indicates that HAWKS was able to produce datasets that single-linkage performed both very well and very poorly on, dependent on whether it was designated as the winner or loser of the head-to-head.

There are some clear differences between certain pairings of algorithms. As hinted in Section V-A, HAWKS is consistently able to generate maximal (i.e. $\phi(\mathcal{A}_w) - \phi(\mathcal{A}_l) \approx 1$) performance difference when single-linkage is \mathcal{A}_l , but the performance difference appears minimal for single-linkage vs. average-linkage, indicating that any weaknesses specific to average-link are not weaknesses that single-link can exploit. The average ARI for average-linkage, GMM and K-Means++ when set as the ‘loser’ (\mathcal{A}_l) indicates a difficulty in consistently generating datasets that these algorithms perform very poorly on. This is not unexpected, given the generating model used. As long as clusters are not fully overlapping (which we discourage with our *overlap* constraint) we expect these algorithms to identify some elements of the clusters, making an ARI close to 0 unlikely.

To evaluate the influence of initialization on these results, in Fig. 5b the stochastic algorithms (GMM and K-Means++) are re-run with a different initialization. While K-Means++ was largely unchanged, there was an average ARI increase (as \mathcal{A}_l) of 0.13 for GMM, highlighting that in some cases the lower performance was due purely to a poor initialization. The average ARI did, however, slightly decrease (as \mathcal{A}_w), highlighting the potential disadvantage of these stochastic algorithms over the linkage-based algorithms as they converge to local minima. This suggests another potential use case for HAWKS. When given an infinite budget of initializations and picking the best, we expect these algorithms to perform better. The robustness of the initialization method can be assessed, however, by investigating how often the algorithm is still able to achieve good performance with a limited budget.

In order to investigate the ability of the *Versus* mode to grant algorithmic insight, we need to inspect the structures HAWKS discovered for different algorithm combinations. For this, it is important to identify where HAWKS struggles to generate datasets that favour one algorithm over another. We can then try to establish whether this is due to the superiority of one algorithm over another, or the inability of HAWKS to generate structures with properties that *would* differentiate them. For brevity, the following sections discuss some interesting examples for a subset of the scenarios (shown in Fig. 6). Further examples of the algorithm pairings can be found in the supplementary material (Section ??).

1) *GMM vs. single-linkage*: Owing to the Gaussian representation that HAWKS uses, and the known issues of single-linkage, this scenario is expected to provide a large performance difference. Fig. 6a shows that HAWKS achieves this performance difference by exploiting the aforementioned ‘chaining’ effect of single-linkage [53]. Discovering this requires iterative movement of the clusters in order for the data points to be close enough to induce this effect, supporting the utility of our mutation operators’ nuanced ability to adjust the location of clusters.

2) *Single-linkage vs. GMM*: The reverse scenario should be much harder for HAWKS, as GMM is largely insensitive to eccentricity and naturally fits our cluster representation. As exemplified in Fig. 6b (and observed in the other datasets produced), HAWKS tends to place large clusters far away from several smaller compact clusters in order to increase the chance of a poor initialization from GMM, exploiting the stochasticity of this method.⁹

3) *GMM vs. K-Means++*: In Fig. 6c we can see that a key exploit, as discovered by HAWKS, is the inability of K-Means++ to handle eccentric clusters. The example highlights significant differences between the use of mixture models and an algorithm relying on assignment to the closest centroid. A clear performance differential is found on this dataset, despite basic similarities in the inductive biases of the two algorithms.

⁹Each fitness evaluation uses a different initialization for GMM and K-Means++, otherwise HAWKS tacitly exploits this knowledge by moving the clusters into the static initial centroid locations. This will maximize performance difference, but represents a technological artefact rather than a generic property.

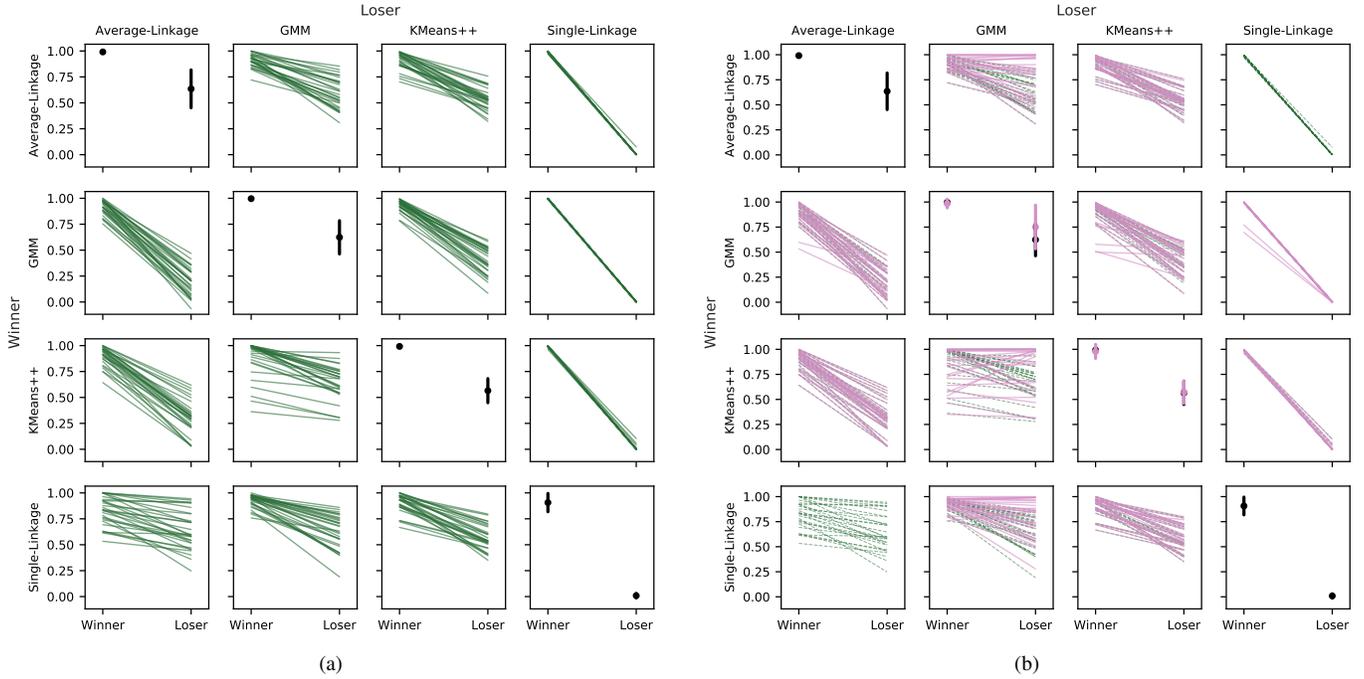


Fig. 5. Performance (ARI) for each algorithm as the winner (row) and loser (column). Each line represents the best individual from a single run, connecting the ARI obtained for \mathcal{A}_w and \mathcal{A}_l . On the diagonal, the average (and standard deviation) performance is aggregated for that individual as the winner and loser, indicating the overall capability of HAWKS to produce datasets that are simple and difficult for that algorithm. The stochastic algorithms (GMM and K-Means++) are re-run with a different initialization (as indicated by the solid lines, with the original results as dashed lines) in (b) to further measure robustness.

4) *K-Means++ vs. GMM*: As both algorithms are well-suited for compact clusters, eccentricity is not a characteristic that can be used for differentiating performance in this scenario. Here, HAWKS exploits GMM’s previously-mentioned weakness (of sub-dividing a single large cluster) during its initialization stage. K-Means++ is less sensitive to this problem due to its improved initialization routine. This shows the utility of HAWKS in identifying relative strengths and weaknesses of specific algorithms, which could aid in algorithmic development (e.g. when empirically comparing initialization schemes).

The datasets in the *Versus* mode are optimized towards a performance differential between the algorithms, rather than towards a cluster structure specified by the cluster validity index (as done in the *Index* mode). It is therefore of interest to investigate how these datasets compare in terms of problem feature diversity. For this, we add the 30 datasets from each of the 12 head-to-heads to the instance space created in Section V-A (using the same principal components).

In Fig. 7, we have highlighted the datasets from the *Index* and *Versus* modes (with the other datasets in grey for reference). Clearly, we are able to generate datasets that are notably different in terms of their problem features (and thus properties). In particular, there are many datasets in the region where previously only *HK* produced datasets, further highlighting the flexibility of our generating mechanism in covering additional regions, as objective function, clustering algorithms, and other parameters are varied.

VI. CONCLUSION

Clustering is a vital tool for pattern discovery, but it is often unclear which clustering algorithm is the most appropriate for a given dataset. An optimal choice requires an accurate understanding of the data properties as well as the strengths and weaknesses of candidate algorithms. Both types of information are difficult to come by in typical real-world settings.

Synthetic benchmark datasets play an important role in improving our understanding of the former, i.e. to examine the specific strengths and capabilities of a given clustering method. Their specific advantage is the availability of a known generating model, which allows researchers to relate aspects of the true cluster structure to algorithm performance. Unfortunately, available synthetic benchmarks for cluster analysis cover a limited variety of structural aspects, and there are no existing generators that have been designed with a wider flexibility in mind.

Our framework HAWKS employs the power of an EA to better meet the challenges highlighted above, and to generate more diverse collections of benchmarks, in particular. When compared to existing clustering benchmarks, HAWKS is found to generate datasets exhibiting more feature diversity and eliciting more variation in algorithm performance. Future work could improve the ability of the instance space to distinguish algorithmic footprints, by further enriching the set of problem features and improving the projection methodology.

Finally, HAWKS can be modified to directly generate datasets that are either simple or difficult for a given algorithm, facilitating a deeper understanding of existing algorithms and

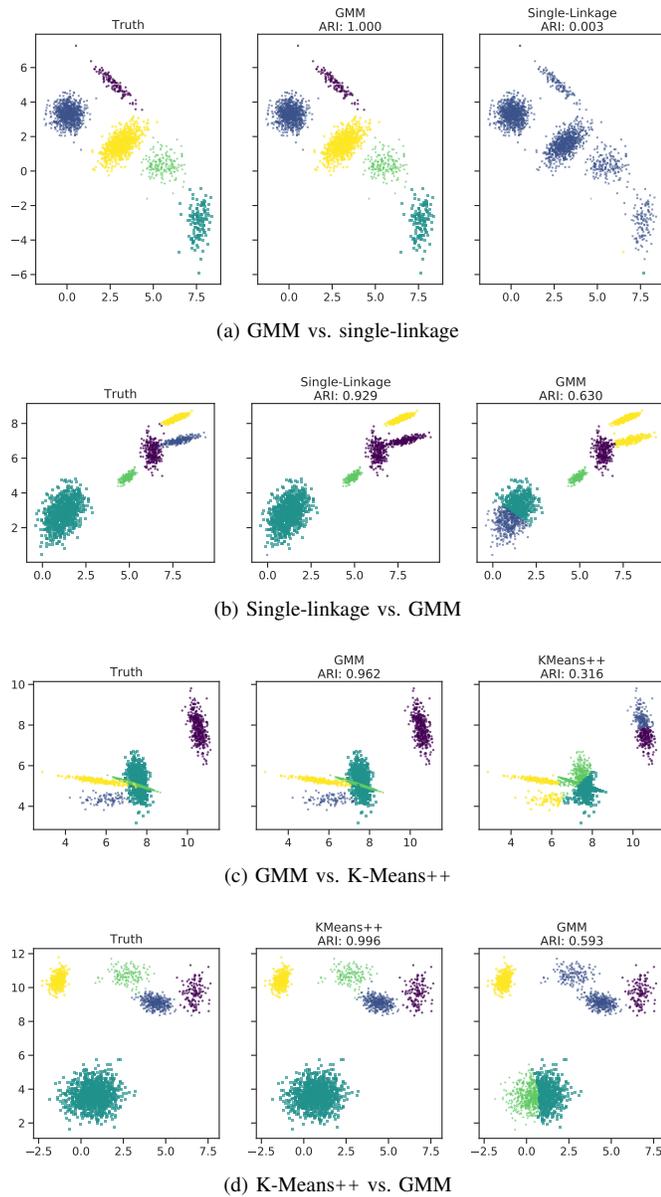


Fig. 6. Examples of datasets for the listed head-to-heads. Each figure shows the ground truth (left column), and the cluster assignment for each algorithm with the associated ARI (middle and right columns).

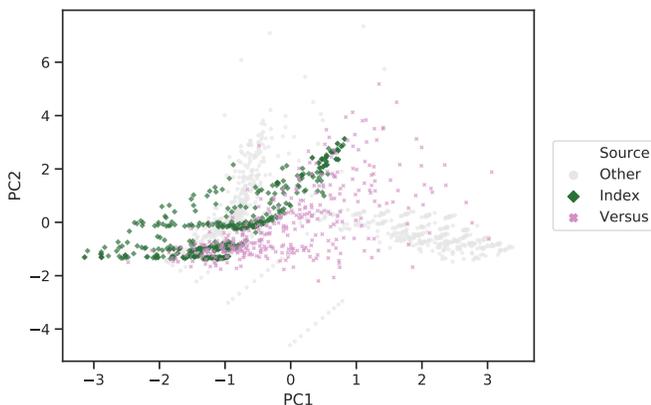


Fig. 7. Instance space created in Section V-A with the datasets produced by the two modes of HAWKS highlighted (the ‘Other’ points are the other datasets).

potentially informing algorithm development. This provides a new avenue for “controlled experimentation” with clustering algorithms, which does not rely on the use of over-simplified toy datasets. Investigations into the use of this *Versus* mode in conjunction with more complex clustering algorithms, could further test HAWKS’ capabilities and may provide novel insights into the inductive biases and performance of these algorithms.

REFERENCES

- [1] J. Handl, J. D. Knowles, and D. B. Kell, “Computational cluster validation in post-genomic data analysis,” *Bioinform.*, vol. 21, no. 15, pp. 3201–3212, 2005.
- [2] U. Maulik, S. Bandyopadhyay, and A. Mukhopadhyay, *Multiobjective Genetic Algorithms for Clustering - Applications in Data Mining and Bioinformatics*. Springer, 2011.
- [3] R. Xu and D. C. Wunsch, “Clustering algorithms in biomedical research: A review,” *IEEE Reviews in Biomedical Engineering*, vol. 3, pp. 120–154, 2010.
- [4] M. Ahmed, A. N. Mahmood], and J. Hu, “A survey of network anomaly detection techniques,” *Journal of Network and Computer Applications*, vol. 60, pp. 19–31, 2016.
- [5] Z. Ma, J. M. R. Tavares, R. N. Jorge, and T. Mascarenhas, “A review of algorithms for medical image segmentation and their applications to the female pelvic cavity,” *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 13, no. 2, pp. 235–246, 2010.
- [6] C. Hennig, “What are the true clusters?” *Pattern Recognit. Lett.*, vol. 64, pp. 53–62, 2015.
- [7] A. K. Jain, M. N. Murty, and P. J. Flynn, “Data clustering: A review,” *ACM Comput. Surv.*, vol. 31, no. 3, pp. 264–323, 1999.
- [8] U. von Luxburg, R. C. Williamson, and I. Guyon, “Clustering: Science or art?” in *Unsupervised and Transfer Learning – Workshop held at ICML 2011*, ser. JMLR Proceedings, vol. 27. JMLR.org, 2012, pp. 65–80. [Online]. Available: <http://proceedings.mlr.press/v27/luxburg12a.html>
- [9] O. Arbelaitz, I. Gurrutxaga, J. Muguerza, J. M. Pérez, and I. Perona, “An extensive comparative study of cluster validity indices,” *Pattern Recognit.*, vol. 46, no. 1, pp. 243–256, 2013.
- [10] J. N. Hooker, “Testing heuristics: We have it all wrong,” *Journal of Heuristics*, vol. 1, no. 1, pp. 33–42, 1995.
- [11] J. Handl and J. D. Knowles, “Improvements to the scalability of multiobjective clustering,” in *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2005, 2-4 September 2005, Edinburgh, UK*. IEEE, 2005, pp. 2372–2379.
- [12] K. Smith-Miles and T. T. Tan, “Measuring algorithm footprints in instance space,” in *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2012*. IEEE, 2012, pp. 1–8.
- [13] N. Macià and E. Bernadó-Mansilla, “Towards uci+: a mindful repository design,” *Information Sciences*, vol. 261, pp. 237–262, 2014.
- [14] O. Mersmann, M. Preuss, and H. Trautmann, “Benchmarking evolutionary algorithms: Towards exploratory landscape analysis,” in *International Conference on Parallel Problem Solving from Nature 2010*. Springer, 2010, pp. 73–82.
- [15] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural Networks*, vol. 61, pp. 85–117, 2015.
- [16] D. H. Wolpert, “The lack of A priori distinctions between learning algorithms,” *Neural Computation*, vol. 8, no. 7, pp. 1341–1390, 1996.
- [17] J. R. Rice, “The algorithm selection problem,” 1976, vol. 15, pp. 65–118.
- [18] K. Smith-Miles, “Cross-disciplinary perspectives on meta-learning for algorithm selection,” *ACM Comput. Surv.*, vol. 41, no. 1, pp. 6:1–6:25, 2008.
- [19] K. Smith-Miles, D. Baatar, B. Wreford, and R. Lewis, “Towards objective measures of algorithm performance across instance space,” *Computers & Operations Research*, vol. 45, pp. 12–24, 2014.
- [20] K. Smith-Miles and L. Lopes, “Measuring instance difficulty for combinatorial optimization problems,” *Computers & Operations Research*, vol. 39, no. 5, pp. 875–889, 2012.
- [21] K. Smith-Miles and S. Bowly, “Generating new test instances by evolving in instance space,” *Computers & Operations Research*, vol. 63, pp. 102–113, 2015.
- [22] M. A. Muñoz, L. Villanova, D. Baatar, and K. Smith-Miles, “Instance spaces for machine learning classification,” *Machine Learning*, vol. 107, no. 1, pp. 109–147, 2018.

- [23] M. Birattari, *Tuning Metaheuristics - A Machine Learning Perspective*, ser. Studies in Computational Intelligence. Springer, 2009, vol. 197.
- [24] F. Hutter, L. Kotthoff, and J. Vanschoren, Eds., *Automated Machine Learning - Methods, Systems, Challenges*, ser. The Springer Series on Challenges in Machine Learning. Springer, 2019.
- [25] M. López-Ibáñez, J. Dubois-Lacoste, L. P. Cáceres, M. Birattari, and T. Stütze, “The irace package: Iterated racing for automatic algorithm configuration,” *Operations Research Perspectives*, vol. 3, pp. 43–58, 2016.
- [26] P. Kerschke and M. Preuss, “Exploratory landscape analysis: advanced tutorial at GECCO 2017,” in *Genetic and Evolutionary Computation Conference, 2017, Companion Material Proceedings*, P. A. N. Bosman, Ed. ACM, 2017, pp. 762–781.
- [27] L. Kotthoff, P. Kerschke, H. H. Hoos, and H. Trautmann, “Improving the state of the art in inexact TSP solving using per-instance algorithm selection,” in *Learning and Intelligent Optimization - 9th International Conference, LION 2015. Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 8994. Springer, 2015, pp. 202–217.
- [28] O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, and G. Rudolph, “Exploratory landscape analysis,” in *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2011*, N. Krasnogor and P. L. Lanzi, Eds. ACM, 2011, pp. 829–836.
- [29] G. C. Bowker and S. L. Star, *Sorting Things Out: Classification and its Consequences*. MIT Press, 2000.
- [30] B. S. Everitt, S. Landau, M. Leese, and D. Stahl, *Cluster Analysis*, 5th ed. Wiley Publishing, 2011.
- [31] S. Ben-David, “Clustering — what both theoreticians and practitioners are doing wrong,” in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence 2018*. AAAI Press, 2018, pp. 7962–7964. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17420>
- [32] J. M. Kleinberg, “An impossibility theorem for clustering,” in *Advances in neural information processing systems*, 2003, pp. 463–470.
- [33] D. G. Ferrari and L. N. de Castro, “Clustering algorithm selection by meta-learning systems: A new distance-based problem characterization and ranking combination methods,” *Inf. Sci.*, vol. 301, pp. 181–194, 2015.
- [34] R. G. F. Soares, T. B. Ludermir, and F. de A. T. de Carvalho, “An analysis of meta-learning techniques for ranking clustering algorithms applied to artificial data,” in *Artificial Neural Networks - ICANN 2009*, ser. Lecture Notes in Computer Science, vol. 5768. Springer, 2009, pp. 131–140.
- [35] D. Dheeru and E. Karra Taniskidou, “UCI machine learning repository,” 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [36] P. Fränti and S. Sieranoja, “K-means properties on six clustering benchmark datasets,” *Appl. Intell.*, vol. 48, no. 12, pp. 4743–4759, 2018.
- [37] J. Handl and J. Knowles, “Feature subset selection in unsupervised learning via multiobjective optimization,” *International Journal of Computational Intelligence Research*, vol. 2, no. 3, pp. 217–238, 2006.
- [38] W. Qiu and H. Joe, “Generation of random clusters with specified degree of separation,” *J. Classification*, vol. 23, no. 2, pp. 315–334, 2006.
- [39] —, “Separation index and partial membership for clustering,” *Computational statistics & data analysis*, vol. 50, no. 3, pp. 585–603, 2006.
- [40] C. Shand, R. Allmendinger, J. Handl, A. M. Webb, and J. Keane, “Evolving controllably difficult datasets for clustering,” in *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2019*. ACM, 2019, pp. 463–471.
- [41] G. W. Stewart, “The efficient generation of random orthogonal matrices with an application to condition estimators,” *SIAM Journal on Numerical Analysis*, vol. 17, no. 3, pp. 403–409, 1980. [Online]. Available: <http://www.jstor.org/stable/2156882>
- [42] S. E. Akhanli and C. Hennig, “Comparing clusterings and numbers of clusters by aggregation of calibrated clustering validity indexes,” *arXiv preprint arXiv:2002.01822*, 2020.
- [43] P. J. Rousseeuw, “Silhouettes: a graphical aid to the interpretation and validation of cluster analysis,” *Journal of computational and applied mathematics*, vol. 20, pp. 53–65, 1987.
- [44] J. Xie, R. Girshick, and A. Farhadi, “Unsupervised deep embedding for clustering analysis,” in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ser. ICML’16. JMLR.org, 2016, pp. 478–487.
- [45] L. Hubert and P. Arabie, “Comparing partitions,” *J. Classification*, vol. 2, no. 1, pp. 193–218, 1985.
- [46] T. P. Runarsson and X. Yao, “Stochastic ranking for constrained evolutionary optimization,” *IEEE Trans. on Evolutionary Computation*, vol. 4, no. 3, pp. 284–294, 9 2000.
- [47] J. Kennedy and R. Eberhart, “Particle swarm optimization,” in *Proceedings of International Conference on Neural Networks (ICNN’95)*. IEEE, 1995, pp. 1942–1948.
- [48] B. Li, K. Tang, J. Li, and X. Yao, “Stochastic ranking algorithm for many-objective optimization based on multiple indicators,” *IEEE Trans. Evolutionary Computation*, vol. 20, no. 6, pp. 924–938, 2016.
- [49] J. Handl and J. D. Knowles, “An Evolutionary Approach to Multiobjective Clustering,” *IEEE Trans. on Evolutionary Computation*, vol. 11, no. 1, pp. 56–76, 2007.
- [50] C. H. Q. Ding and X. He, “K-nearest-neighbor consistency in data clustering: incorporating local information into global optimization,” in *Proceedings of the 2004 ACM Symposium on Applied Computing (SAC)*, H. Haddad, A. Omicini, R. L. Wainwright, and L. M. Liebrock, Eds. ACM, 2004, pp. 584–589.
- [51] M. Garza-Fabre, J. Handl, and J. D. Knowles, “An Improved and More Scalable Evolutionary Approach to Multiobjective Clustering,” *IEEE Trans. on Evolutionary Computation*, vol. 22, no. 4, pp. 515–535, 2017.
- [52] D. Arthur and S. Vassilvitskii, “k-means++: The advantages of careful seeding,” in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007*. SIAM, 2007, pp. 1027–1035. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1283383.1283494>
- [53] L. Hubert, “Approximate evaluation techniques for the single-link and complete-link hierarchical clustering procedures,” *Journal of the American Statistical Association*, vol. 69, no. 347, pp. 698–704, 1974.
- [54] J. Demšar, “Statistical comparisons of classifiers over multiple data sets,” *Journal of Machine Learning Research*, vol. 7, no. Jan, pp. 1–30, 2006. [Online]. Available: <http://jmlr.org/papers/v7/demars06a.html>
- [55] P. Nemenyi, “Distribution-free multiple comparisons,” Ph.D. dissertation, Princeton University, 1963.

HAWKS: Evolving Challenging Benchmark Sets for Cluster Analysis

– SUPPLEMENTARY MATERIAL –

Cameron Shand, Richard Allmendinger, Julia Handl, Andrew Webb, and John Keane

S-I. HAWKS DETAILS

A. Generating random cluster sizes to a pre-defined limit

To generate randomly sized clusters that sum to a given value with a minimum size, the following approach is used:

$$\mathbf{w} = [w_1, \dots, w_K]$$

where K is the number of clusters, $w_i \sim \text{Dir}(\alpha)$, and $\alpha = (1, 1, \dots, 1)$. For a given cluster i , the size is:

$$|C_i| = C_{\min} + w_i(N - (K \times C_{\min}))$$

where C_{\min} is the minimum size a cluster can be.

This ensures a uniform distribution of cluster sizes *after* scaling such that they sum to N , which is not guaranteed when simply sampling random numbers and scaling these to N (as this no longer guarantees a uniform distribution).

B. Scaling the covariance matrix

We provide further details about the covariance mutation operator described in Section ???. To mutate the covariance matrix in HAWKS, two separate matrices are generated for the perturbation. The rotation matrix randomly rotates the cluster, whereas the scaling matrix modifies the magnitude of the principal semi-axes, modifying the area of the space that is covered by this distribution.

To avoid issues of converging towards either spherical or highly eccentric clusters upon multiple applications of the mutation operator, we ensure that the determinant of the covariance matrix remains the same after applying the scaling matrix. For this, let x_i be the elements of a vector sampled from a Dirichlet distribution. Then

$$\begin{aligned} \sum_i^D x_i &= 1 \\ \Rightarrow \sum_i^D \left(x_i - \frac{1}{D}\right) &= 0 \\ \Rightarrow \exp\left(\sum_i^D \left(x_i - \frac{1}{D}\right)\right) &= 1 \\ \Rightarrow \prod_i^D \exp\left(x_i - \frac{1}{D}\right) &= 1. \end{aligned}$$

The determinant of a diagonal matrix is the product of the values on the diagonal, thus the scaling matrix $(\exp(x_i - \frac{1}{D}))$ has determinant 1.

C. Average eccentricity problem feature

In order to reduce the sensitivity to outliers and subspace clusters when calculating the average eccentricity, we use a subset of the eigenvalues (obtained via singular value decomposition) that account for 95% of the total sum of the eigenvalues. This is functionally identical to using the principal components (obtained via principal component analysis) that account for 95% of the variance.

S-II. MUTATING CLUSTERS IN HIGHER DIMENSIONS

Shand et al. [1] noted that the original mutation operator for the cluster means (μ) is decreasingly useful as the dimensionality increases due to stochastic movement of the mean, as there are an increasing number of directions to move away from other clusters. This results in a bias towards increasing the silhouette width as the clusters drift apart. In this section, we propose and compare several different operators that try to explicitly incorporate directionality into the random movement to avoid this bias.

Incorporating explicit directionality into the operator can directly guide whether the movement of clusters is either away from or towards existing clusters, increasing and decreasing the silhouette width respectively. For this, the operator would need to incorporate the position of at least one other cluster in the random perturbation.

The previous [1] mutation operator was defined as:

$$\mu'^{(i)} \sim \mathcal{N}(\mu^{(i)}, s) \quad (\text{S-1})$$

where $\mu'^{(i)}$ is the new mean, $\mu^{(i)}$ is the current mean, and s is the width (variance) of the multivariate normal distribution. In other words, the new mean is sampled from a normal distribution around the current mean. Next, we define each of the proposed operators.

A. Mutation Operator Descriptions

1) “*Rails*” operator: As a simple baseline to examine the utility of including directionality, this operator randomly selects a cluster, and the current cluster moves either away from or towards that cluster with a random weighting ($0 \leq w_1 \leq 1$).

To select a random cluster that is different from the current one, let n be a random integer from the set $\{1, \dots, K\} \setminus \{i\}$. The mean of the current cluster, $\mu^{(i)}$, can be mutated as follows:

$$\mu'^{(i)} = \begin{cases} \mu^{(i)} + [w_1(\mu^{(n)} - \mu^{(i)})] & \text{if } p \leq 0.5 \\ \mu^{(i)} - [w_1(\mu^{(n)} - \mu^{(i)})] & \text{if } p > 0.5. \end{cases} \quad (\text{S-2})$$

where $\mu^{(n)}$ is the mean of the random cluster n , $w_1 \sim U(0, 1)$ is a random weight, $\mu^{(n)} - \mu^{(i)}$ is the difference between the cluster means, and $p \sim U(0, 1)$ is a random uniform probability to determine whether we mutate towards or away from the other cluster.

2) *PSO-inspired mutation with random directionality*: For a mutation that covers more of the space, rather than along the vectors between centroids, we take inspiration from another EA paradigm: particle swarm optimization (PSO). In the original PSO mutation operator, each particle is updated by incorporating their current position, the best position that particle has ever had, and the best position ever found by any particle. These latter two best positions are weighted by random coefficients in order to create a random movement of the particles. For further details, see [2], [3]. We use this as inspiration by viewing the clusters as particles, and try to mutate the location of the cluster based on the position of another cluster and some global representative. As our fitness is derived from a combination of all particles, the original notions of personal and global best are not applicable here, but the concept of independently weighting both a single point and an aggregated one is.

By updating the mean using a randomly weighted combination of an existing cluster and the global mean ($\bar{\mu}$) of all clusters, we can create a random movement of the cluster that still takes into account the position of the existing clusters. By incorporating the global mean, we can avoid generating well-separated groups of clusters that can deceive the silhouette width. To ensure that the location of the global mean is calculated in a way meaningful to the fitness, we calculate the mean across all *data points* (not across cluster means), as it is the former that is used in the calculation of the silhouette width (and thus fitness). As a result, relative cluster sizes are incorporated into the mutation. The mean of the current cluster, $\mu^{(i)}$, is thus mutated using the follow operator (referred to as PSO-random):

$$\mu'^{(i)} = \begin{cases} \mu^{(i)} + [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } p \leq 0.5 \\ \mu^{(i)} - [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } p > 0.5. \end{cases} \quad (\text{S-3})$$

3) *PSO-inspired mutation with informed directionality*: By using a coin-flip to determine whether the cluster should move towards or away from existing clusters, there arises obvious scenarios where the existing clusters may be very far apart and a closer structure is desired. An unfavourable coin-flip moves this cluster further away, wasting a perturbation and thus evaluation. By using the sign of the difference between the current silhouette width of the individual and the target ($s_{all} - s_t$), we can move the cluster centre in the direction that will have the best chance of improving the fitness.

The mean of the current cluster, $\mu^{(i)}$, can be mutated as follows:

$$\mu'^{(i)} = \begin{cases} \mu^{(i)} + [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } s_{all} > s_t \\ \mu^{(i)} - [w_1(\mu^{(n)} - \mu^{(i)}) + w_2(\bar{\mu} - \mu^{(i)})] & \text{if } s_{all} \leq s_t. \end{cases} \quad (\text{S-4})$$

The addition of this information will improve convergence in situations where $s_{all} - s_t$ is large, as the fitness will be improved more quickly, but this adds in a bias towards the

fitness (as opposed to the constraints), which is typically controlled through the stochastic ranking (P_f). This operator is referred to as PSO-informed.

4) *DE-inspired mutation*: Taking inspiration from yet another EA paradigm, differential evolution (DE), we view the existing clusters as individual vectors which can help in the creation of a “donor vector” (the new mean) from a “target vector” (the current mean). The classical DE mutation operator combines three existing individuals in order to create a new individual [4]–[6]. We adapt this to generate a new cluster mean from existing ones as follows:

$$\mu'^{(i)} = \mu^{(i)} + F(\mu^{(r_1)} - \mu^{(r_2)}) \quad (\text{S-5})$$

where i , r_1 , and r_2 are distinct indices of a cluster, and F is a constant factor in the range $[0, 2]$. Unlike with the previous operators, the randomness occurs only in terms of which other means are selected, such that the movement vector for the current mean is a fixed multiple (F) of the vector between the randomly selected means. As a result, the number of possible locations that a cluster can mutate to is finite and related to the number of clusters. This lack of flexibility may either slow convergence when larger jumps are needed, or nuance when small movements to the cluster are needed in the final generations (which is especially useful for adjusting the overlap between clusters).

B. Experimental Setup

To investigate the convergence properties of these operators, and whether they show a bias towards separating or bringing together clusters, we set up the following experiments. We set a low (0.2) and high (0.9) s_t to optimize towards poorly- and well-separated clusters respectively, as well as initializing the clusters either together or far apart. These four scenarios test the general ability of the operators to move cluster centroids, and are tested using $D \in \{2, 50\}$ dimensions to check robustness at a higher dimensionality. For each scenario, HAWKS is run for 30 times to assess the robustness over different initializations, and for 100 generations to assess the stability of the operators. For the DE-inspired operator, we use $F = 1$ to strike a balance between convergence speed and potential oscillation when clusters are close together.

C. Results

Fig. S-1 shows convergence plots for the four different scenarios in $2D$, showing the average silhouette width (and not the fitness, to make it clear when s_{all} is above or below the target) across the generations. The target, s_t , is shown by the dashed horizontal line.

In Fig. S-1a, all mutation operators are quickly able to move the clusters apart to increase the silhouette width, but we see a difference in the stability after this as some operators (predominantly DE and PSO-random) further increase the silhouette width above s_t , highlighting a mechanistic bias. Fig. S-1b has the same low s_t , but the initial clusters are further apart (requiring the optimization to bring clusters together). The DE-inspired operator is unable to decrease the silhouette width, and the optimization is driven by minimizing the *overlap*

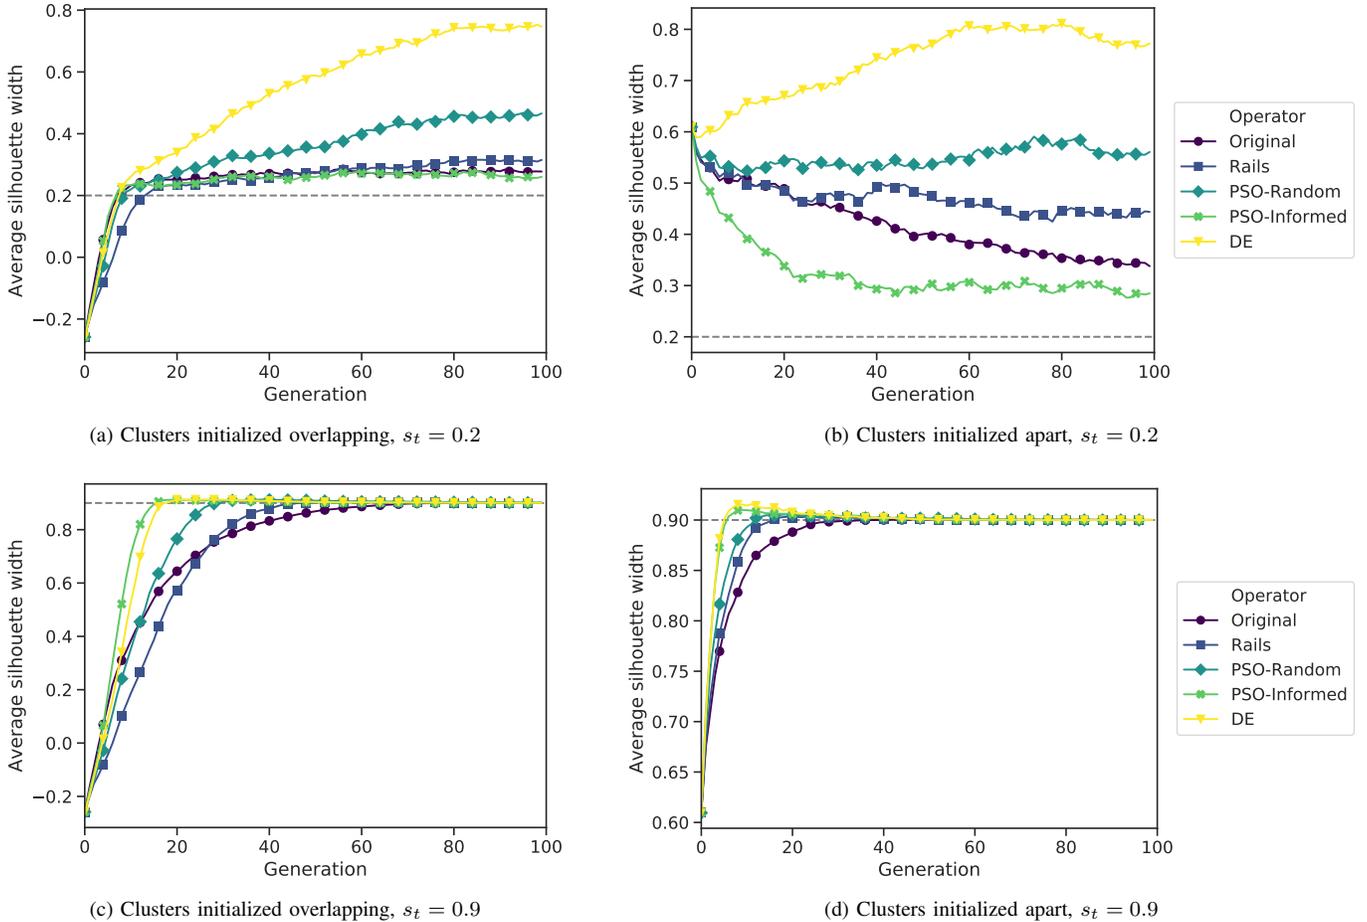


Fig. S-1. Convergence plot showing the average (across all 30 runs) silhouette width across the generations when generating 2-dimensional datasets for our four scenarios. The dashed line shows the target silhouette width (s_t).

constraint (Fig. S-2b). While the PSO-random operator does not converge, the quick decrease in the silhouette width for the PSO-informed operator shows that this is not a mechanistic issue, but one of directionality. Mutually exclusive satisfaction of the fitness (s_t) and the *overlap* constraint prevent any operator from fully reaching $s_t = 0.2$. The slow silhouette width decrease with the original operator highlights the issue with a static step size.

In Fig. S-1c–d we can see that there is no significant difference beyond the speed of convergence between the operators when optimizing for $s_t = 0.9$, independent of the initialization used. No such difficulties with any of the operators can be seen for either initialization in Fig. S-1c–d, where the target silhouette width is high. None of the operators drift once converged, as this target also satisfies the *overlap* constraint.

In Fig. S-2a, we can see the minimization of the *overlap* constraint, which is naturally greater for the operators that increased the silhouette width. This indicates a difference in the relative ease of satisfying s_t and minimizing overlap between the operators, particularly with the lack of pressure (afforded by $P_f = 0.5$) towards either one. With such a low s_t , there is a strong inverse correlation between the overlap and fitness. The PSO-informed operator has an explicit drive to avoid a drift of the silhouette width away from the target,

but the vast difference in drift for the DE-inspired operator illustrates a clear behavioural difference. In Fig. S-2b, when the clusters are initialized apart (and thus the task is to bring them closer together), we can see that the DE-inspired does consistently minimize the *overlap* (as it erroneously increases the silhouette width as the clusters drift apart). The other operators also decrease the *overlap*, apart from PSO-informed where the embedded preference towards the fitness disturbs the minimization of the *overlap*.

This preference is further illustrated when we look at Fig. S-3, where the silhouette width and *overlap* of the best individuals from each run are plotted when the target silhouette width is low ($s_t = 0.2$). The PSO-informed operator creates far less diversity in terms of the *overlap*, thus generating less diverse datasets. To further investigate these operators, we now look at the same scenarios in 50D.

Fig. S-4 shows convergence plots for the four different scenarios in 50D, which was previously where issues with the original operator were first identified. In Fig. S-4a, when the scenario is to keep clusters close together, the drift of the DE-inspired operator is more pronounced relative to the other operators, which remain mostly stable. The bias of the DE-inspired operator towards moving clusters further apart is further shown in Fig. S-4b, where the original operator

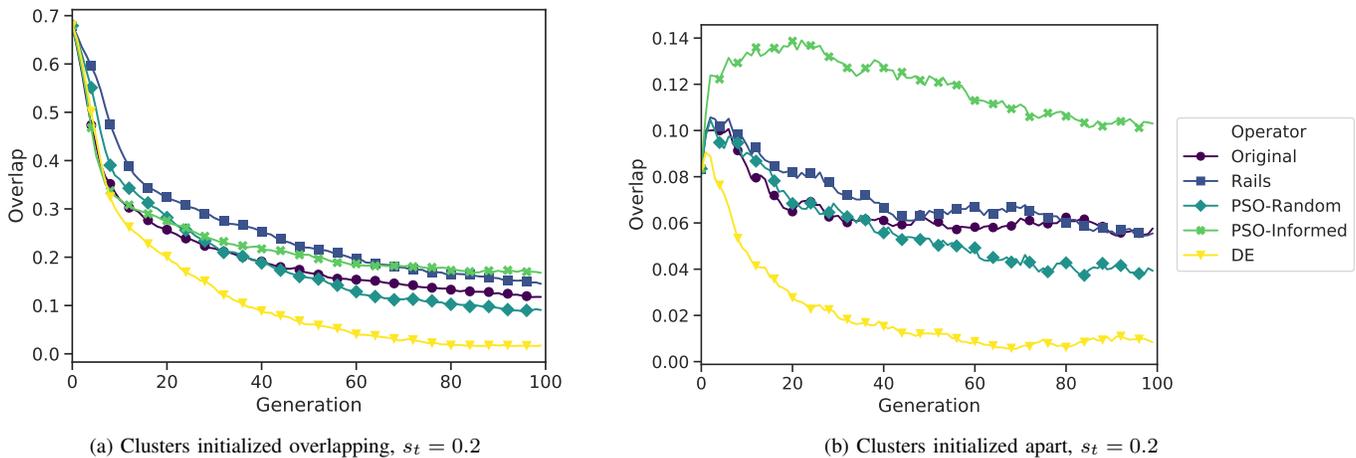


Fig. S-2. Convergence plot showing the average (across all 30 runs) of the *overlap* constraint across the generations when generating 2-dimensional datasets.

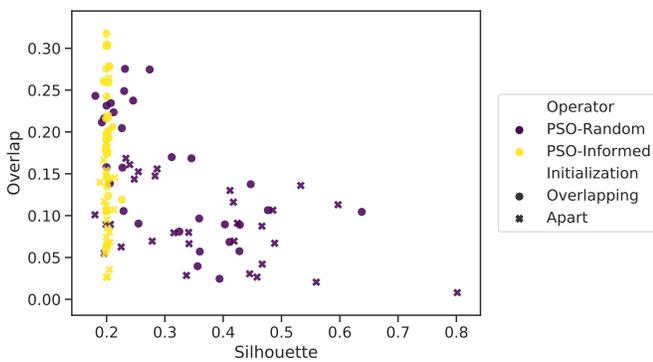


Fig. S-3. Scatter plot of the silhouette width and *overlap* constraint for each of the best individuals from every run for the low silhouette width target.

is also incapable of bringing clusters closer together. As the DE-inspired operator uses the direction and magnitude of the vector between two random clusters, it does not necessarily move the cluster being mutated in the direction of other clusters, an issue which is more prevalent in higher dimensions. As expected, the PSO-informed operator converges the fastest, whereas the “rails” and PSO-random operators increasingly slow to decrease the silhouette width (largely due to trying to minimize the overlap).

The ability of our original operator to move clusters away from each other in higher dimensions is highlighted in Fig. S-4c, which is clearly contrasted with the proposed operators that are all able to converge significantly faster. Once again, utilizing the individual’s current silhouette width allows the PSO-informed operator to consistently move the clusters apart, converging rapidly. The speed of convergence for the DE-inspired operator is likely due to its step size being based on the distance between two random clusters, thus enabling it to rapidly move clusters apart. In Fig. S-4d, when the clusters are initialized further apart, similar behaviour to the low s_t is seen. The initial silhouette width is very close to the target, yet our original and the DE-inspired operators are still unable to significantly improve the fitness as they begin *above* the target. It is likely that in this scenario, the step size is too large for

the DE-operator to be useful, and too small and undirected for our original operator.

As such, in our experiments we select to use the PSO-random operator, as it provides a more nuanced mechanism for moving cluster means, but avoids introducing additional bias which would otherwise be controlled by stochastic ranking. For scenarios where the fitness is prioritized, PSO-informed may be the more useful operator, but for general use the PSO-random operator is preferred.

S-III. ANALYZING BENCHMARK DATASET DIVERSITY (INDEX MODE)

In Section ??, we compare several collections of datasets to a set of datasets produced by HAWKS. Here, we provide some further information on the different properties of these datasets, and further supporting results.

A. Dataset collection details

Table S-I shows some basic parameters about the different collections of datasets. Where possible, values are condensed into a range (i.e. “6–11”), but are otherwise listed. Owing to the different nature of these datasets (some real-world, some from a generator), some values are specific to a single dataset, and others represent multiple instantiations from a generator.

B. Problem feature values across the instance space

Fig. ?? showed the instance space created for these datasets using our set of problem features. Fig. S-5 shows this instance space, but with the best ARI achieved by anything algorithm and the values for each of the problem features highlighted.

As we can see, for each problem feature there is a clear gradient across the space, highlighting the utility of using PCA to construct the instance space in terms of understanding how the problem features vary across the space. Interestingly, there is also a gradient for the best ARI achieved, for which lower values are generally associated with either a lower silhouette width, higher standard deviation of the silhouette width, and higher average cluster eccentricity. Overall, this highlights that using problem features specific to clustering help to create

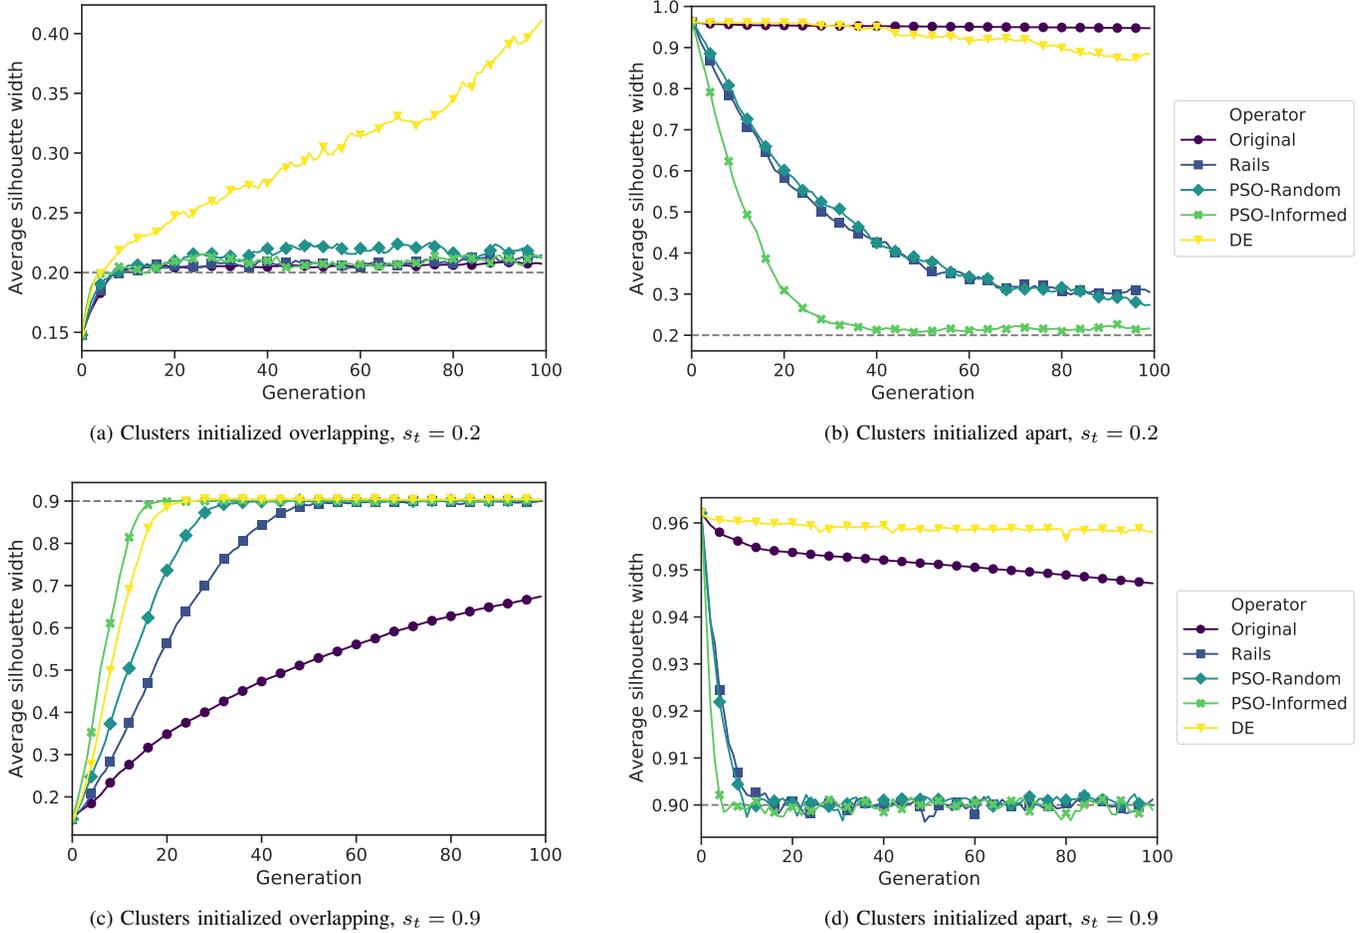


Fig. S-4. Convergence plot showing the average (across all 30 runs) silhouette width across the generations when generating 50-dimensional datasets for our four scenarios.

TABLE S-I
THE NUMBER OF DATA POINTS (N), CLUSTERS (K), DIMENSIONS (D), AND DATASETS FOR EACH COLLECTION OF DATASETS.

Source	N^1	K	D	# Datasets
HAWKS	500, 2500	5, 30	2, 50	448
<i>HK</i> [7]	5474 ± 2137	10, 20, 40, 60, 80, 100, 120	20, 50, 100, 150, 200	350
<i>QJ</i> [8]	599 ± 289	3, 6, 9	5–24	243
<i>SIPU</i> [9]	2248 ± 817	2, 15, 20, 35, 50	$2^1, 2^2, \dots, 2^{10}$	107
<i>UCI</i> [10]	600 ± 571	2, 3, 4, 6, 7, 8, 10, 11, 15	3, 4, 6–11, 13, 18, 19, 22, 30, 34, 44, 60, 90, 166	20
<i>UKC</i> [11]	30685 ± 2178	10, 11, 12	2	8

¹ The average (and standard deviation) are given, except for HAWKS for which the values are ranged.

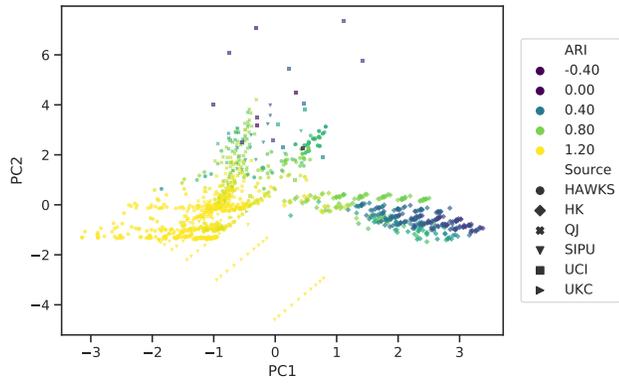
an instance space that correlates to algorithmic performance. As shown in Fig. ?? however, through either the projection method or an incomplete set of problem features (as discussed in Section ??, it is difficult to get a complete set), there is not a complete separation of regions where a single algorithm is superior.

Table S-II shows the mean and standard deviation of each of the problem features for every collection of datasets studied in this paper, clarifying the observations seen for the instance space visualizations (Fig. S-5). The *UCI* datasets show a much higher connectivity, indicating that the nearest neighbours of data points are typically in a different cluster, and thus have poor cluster structure.

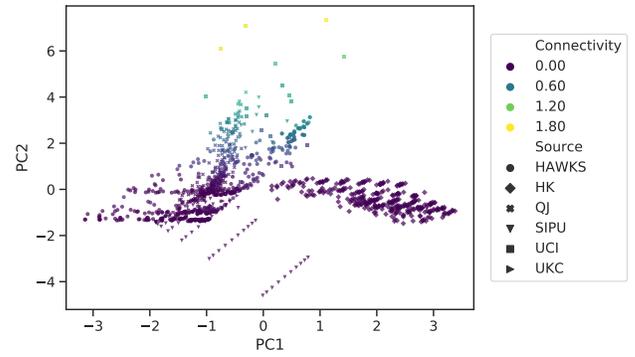
We can compare the mean values for the problem features for the *Index* and *Versus* modes to see if the different optimization settings had a particular bias towards e.g. a low connectivity. The main differences between these datasets, according to the problem features, is that the *Versus* datasets have a slightly higher average connectivity, much lower eccentricity, and higher variance in the silhouette width. The lower eccentricity can be explained through both the initialization used and low dimensionality. As shown in our examples, where highly eccentric clusters were seen, this did not diminish HAWKS' capability of evolving such clusters when needed.

C. Comparing performance diversity via critical difference diagrams

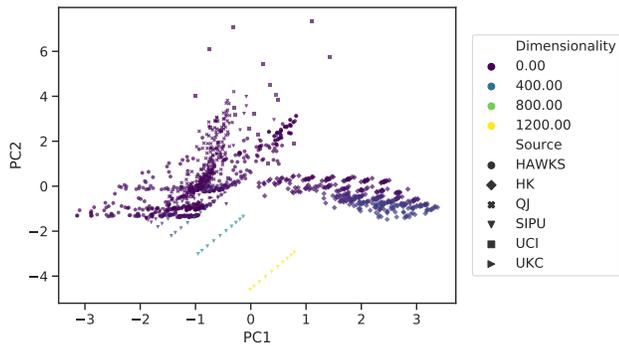
In the main paper, we showed the critical difference (CD) diagrams for datasets from the HAWKS and *HK* generators.



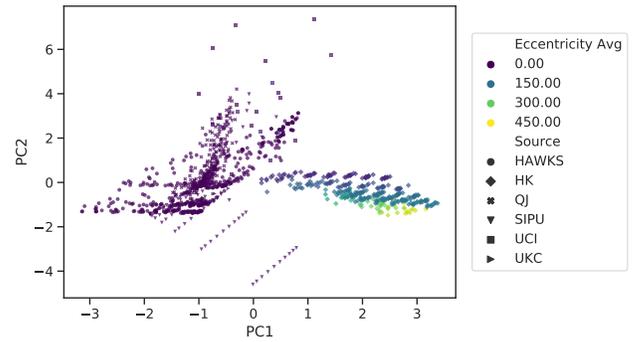
(a) ARI



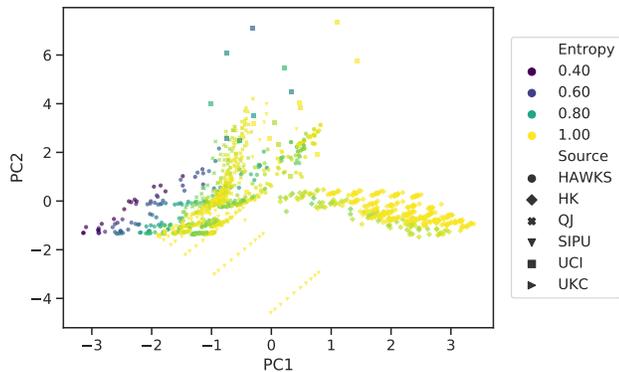
(b) Connectivity



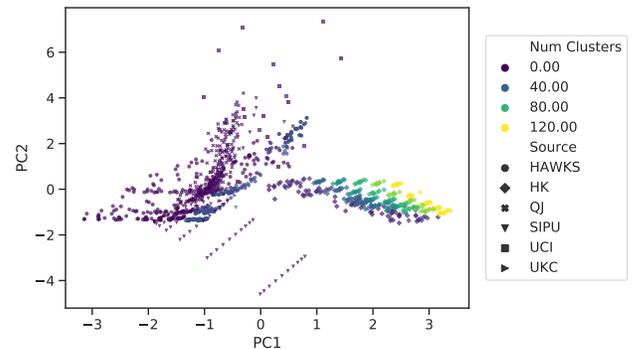
(c) Number of dimensions



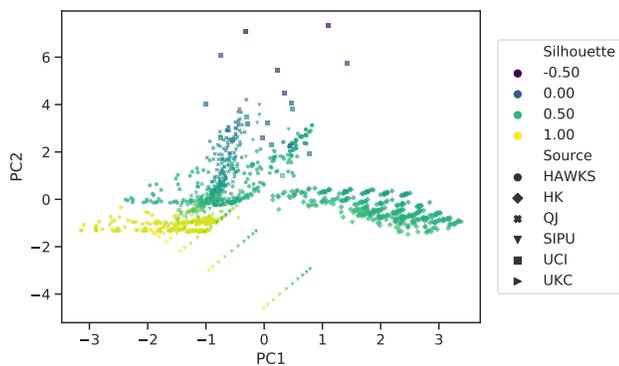
(d) Average eccentricity



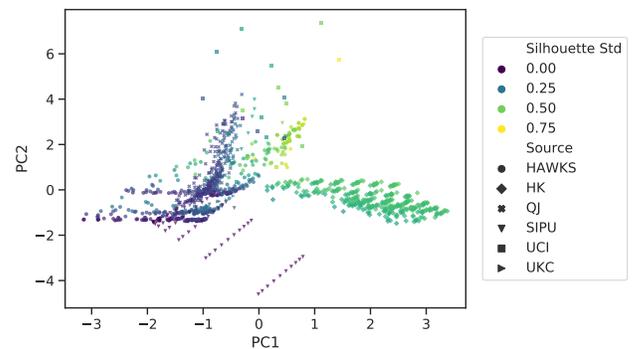
(e) Entropy of cluster sizes



(f) Number of clusters



(g) Silhouette width (average)



(h) Silhouette width (standard deviation)

Fig. S-5. The instance space created in Section ?? showing how the values of each of the 7 problem features and the best ARI found by any algorithm vary across the space.

TABLE S-II
THE MEAN AND STANDARD DEVIATION (SD) OF THE PROBLEM FEATURE VALUES FOR EVERY COLLECTION OF DATASETS, INCLUDING BOTH THE INDEX AND VERSUS MODES OF HAWKS.

Collection	Connectivity		Number of Dimensions		Average Eccentricity		Cluster Size Entropy		Number of Clusters		Silh. width (average)		Silh. width (standard deviation)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
HAWKS (<i>Index</i>) ¹	0.114	0.198	26.000	24.027	5.692	2.568	0.884	0.126	17.500	12.514	0.675	0.225	0.226	0.166
HAWKS (<i>Versus</i>) ²	0.144	0.175	2.000	0.000	2.551	1.309	0.825	0.074	5.000	0.000	0.712	0.188	0.300	0.138
<i>HK</i> [7]	0.008	0.008	104.000	65.393	145.202	79.847	0.984	0.021	61.429	38.012	0.474	0.043	0.437	0.023
<i>QJ</i> [8]	0.207	0.247	13.000	6.926	3.615	0.867	0.966	0.023	6.000	2.455	0.380	0.154	0.145	0.036
<i>SIPU</i> [9]	0.091	0.209	191.346	308.120	3.411	3.008	1.000	0.000	3.411	6.280	0.634	0.225	0.091	0.113
<i>UCI</i> [10]	0.789	0.409	28.350	39.125	20.716	20.706	0.897	0.116	4.900	3.712	0.084	0.279	0.375	0.150
<i>UKC</i> [11]	0.001	0.001	2.000	0.000	1.938	0.177	0.997	0.002	10.875	0.641	0.780	0.062	0.232	0.058

¹ The 448 datasets generated in Section ??.

² The 360 datasets generated in Section ??.

For completeness, here in Fig. 6 we show the CD diagrams for the other dataset collections.

For the *QJ* datasets (Fig. 6a), the compactness-based algorithms (GMM and K-Means++) are clearly the best performing algorithms, with GMM nearly superior across all datasets, highlighting a lack of diversity. Similarly, for the *SIPU* datasets (Fig. 6b) GMM and K-Means++ are more equally the best performing, where the lower connectivity (i.e. lack of overlap between clusters) compared to *QJ* datasets does not distinguish performance further between these two algorithms. The *UCI* datasets (Fig. 6c) show little significant difference between the clustering algorithms which, when combined with the low performance shown in Fig. ??, indicates that there is a high variance of performance, but that this diversity is due to low cluster structure (rather than varied cluster structures). Finally, the *UKC* datasets (Fig. 6d) are too few in number to identify significant differences, though the the higher rank of the compactness-based algorithms fits with visualization of the clusters [11].

S-IV. GENERATING DATASETS THAT CHALLENGE SPECIFIC ALGORITHMS (VERSUS MODE)

Section ?? presented results for the reformulation of the fitness function to evolve datasets that directly maximize the performance difference between two algorithms. Here, we provide further visual examples of the different head-to-head scenarios to illustrate the properties which HAWKS found differentiated the algorithms, and further details for the head-to-head between average- and single-linkage.

A. Additional scenarios

1) *Single-linkage vs. Average-linkage*: Fig. ?? showed that there was a high spread in the ARI for both algorithms, and high variance in the maximum difference found between runs, highlighting the difficulty for this scenario. There are some datasets that correlate both high and low performance between the two, indicating that the algorithms do share some similarities. Fig. S-7a shows an example, where the structure is common among the datasets that did find a performance

difference. The clusters are in general well-separated, as by placing some clusters much further away and co-locating a few smaller clusters, the averaging criterion (to determine where to split next) used by average-linkage assigns clusters that are closer together the same cluster label. The mixed eccentricities of clusters that HAWKS is able to generate helps facilitate the discovery of this exploit. As the clusters need to be sufficiently far away before a chance in the fitness can be found, there may be insufficient pressure for this exploit to be reliably found, highlighting the high variation in performance for this scenario.

2) *Average-linkage vs. Single-linkage*: As average-linkage uses the average of the distances between groups, rather than the smallest distance, it is not susceptible to the chaining effect [12] that single-linkage is. As a result, Fig. S-7b shows that the exploit used in this head-to-head is simply to place clusters close together in order for single-linkage to determine the majority of the points are in a single cluster.

REFERENCES

- [1] C. Shand, R. Allmendinger, J. Handl, A. M. Webb, and J. Keane, "Evolving controllably difficult datasets for clustering," in *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2019*. ACM, 2019, pp. 463–471.
- [2] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of International Conference on Neural Networks (ICNN'95)*. IEEE, 1995, pp. 1942–1948.
- [3] J. Kennedy, "Particle swarm optimization," pp. 760–766, 2010.
- [4] R. Storn and K. Price, "Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces," *Journal of Global Optimization*, vol. 11, no. 4, pp. 341–359, 1997.
- [5] K. Fleetwood, "An introduction to differential evolution," in *Proceedings of Mathematics and Statistics of Complex Systems (MASCOS) One Day Symposium, 26th November, Brisbane, Australia, 2004*, pp. 785–791.
- [6] M. L. Ortiz and N. Xiong, "Investigation of mutation strategies in differential evolution for solving global optimization problems," in *Artificial Intelligence and Soft Computing - 13th International Conference, ICAISC 2014*, ser. Lecture Notes in Computer Science, L. Rutkowski, M. Korytkowski, R. Scherer, R. Tadeusiewicz, L. A. Zadeh, and J. M. Zurada, Eds., vol. 8467. Springer, 2014, pp. 372–383.
- [7] J. Handl and J. D. Knowles, "Improvements to the scalability of multiobjective clustering," in *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2005, 2-4 September 2005, Edinburgh, UK*. IEEE, 2005, pp. 2372–2379.
- [8] W. Qiu and H. Joe, "Generation of random clusters with specified degree of separation," *J. Classification*, vol. 23, no. 2, pp. 315–334, 2006.

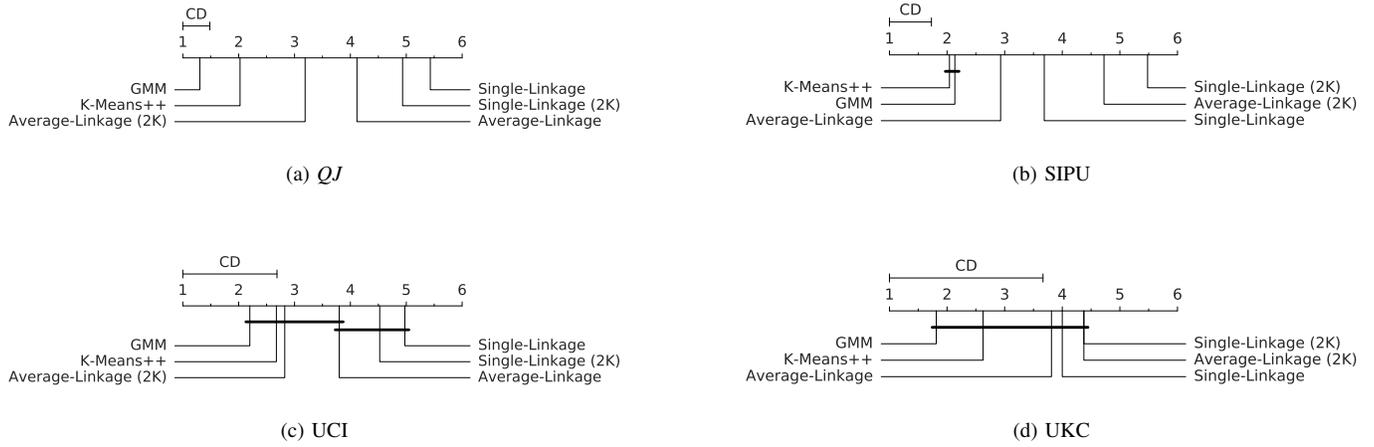


Fig. S-6. Critical difference (CD) diagrams, showing the mean rank (in terms of ARI) for the remaining dataset collections, for each algorithm. Algorithms connected by solid lines are not significantly different according to a two-tailed Nemenyi test.

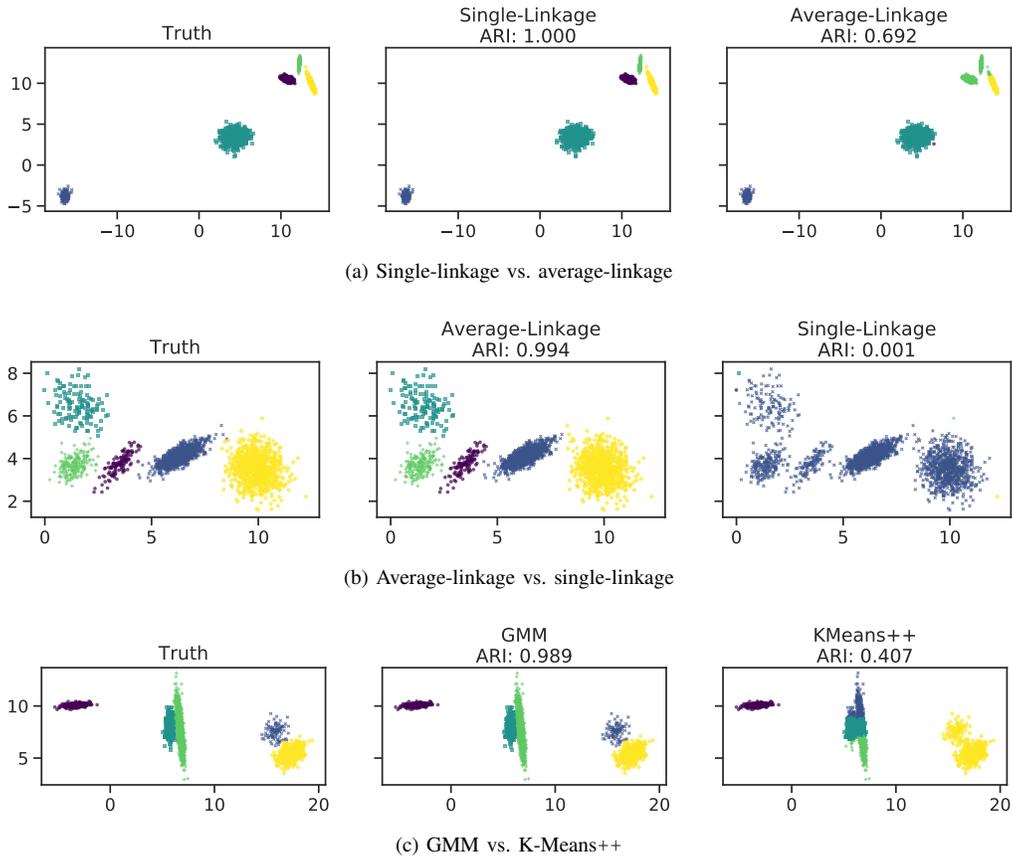


Fig. S-7. Examples of datasets for the listed head-to-heads. Each figure shows the ground truth, and the cluster assignment for each algorithm with the associated ARI.

- [9] P. Fränti and S. Sieranoja, "K-means properties on six clustering benchmark datasets," *Appl. Intell.*, vol. 48, no. 12, pp. 4743–4759, 2018.
- [10] D. Dheeru and E. Karra Taniskidou, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [11] M. Garza-Fabre, J. Handl, and J. D. Knowles, "An Improved and More Scalable Evolutionary Approach to Multiobjective Clustering," *IEEE Trans. on Evolutionary Computation*, vol. 22, no. 4, pp. 515–535, 2017.
- [12] L. Hubert, "Approximate evaluation techniques for the single-link and complete-link hierarchical clustering procedures," *Journal of the American Statistical Association*, vol. 69, no. 347, pp. 698–704, 1974.