Towards Property-Based Classification of Clustering Paradigms

Margareta Ackerman, Shai Ben-David, and David Loker
D.R.C. School of Computer Science
University of Waterloo, Canada

Abstract

Clustering is a basic data mining task with a wide variety of applications. Not surprisingly, there exist many clustering algorithms. However, clustering is an ill defined problem - given a data set, it is not clear what a "correct" clustering for that set is. Indeed, different algorithms may yield dramatically different outputs for the same input sets. Faced with a concrete clustering task, a user needs to choose an appropriate clustering algorithm. Currently, such decisions are often made in a very ad hoc, if not completely random, manner. Given the crucial effect of the choice of a clustering algorithm on the resulting clustering, this state of affairs is truly regrettable. In this paper we address the major research challenge of developing tools for helping users make more informed decisions when they come to pick a clustering tool for their data. This is, of course, a very ambitious endeavor, and in this paper, we make some first steps towards this goal. We propose to address this problem by distilling abstract properties of the input-output behavior of different clustering paradigms.

In this paper, we demonstrate how abstract, intuitive properties of clustering functions can be used to taxonomize a set of popular clustering algorithmic paradigms. On top of addressing deterministic clustering algorithms, we also propose similar properties for randomized algorithms and use them to highlight functional differences between different common implementations of k-means clustering. We also study relationships between the properties, independent of any particular algorithm. In particular, we strengthen Kleinbergs famous impossibility result, while providing a simpler proof.

1 Introduction

In spite of the wide use of clustering in many practical applications, currently, there exists no principled method to guide the selection of a clustering algorithm. Of course, users are aware of the costs involved in employing different clustering algorithms (software purchasing costs, running times, memory requirements, needs for data preprocessing etc.) but there is very little understanding of the differences in the *outcomes* that these algorithms may produce. We focus on that aspect - the input-output properties of different clustering algorithms.

The choice of an appropriate clustering should, of course, be task dependent. A clustering that works well for one task may be unsuitable for another. Even more than for supervised learning, for clustering, the choice of an algorithm must incorporate domain knowledge. While some domain knowledge is embedded in the choice of similarity between domain elements (or the embedding of these elements into some Euclidean space), there is still a large variance in the behavior of difference clustering paradigms over a fixed similarity measure. For some clustering tasks, there is a natural clustering objective function that one may wish to optimize (like k-means for vector quantization coding tasks), but very often the task does not readily translate into a corresponding objective function. Often users are merely searching for a meaningful clustering, without a prior preference for any specific objective function. What seems to be missing is a clear understanding of the differences in clustering outputs in terms of intuitive and usable properties. Furthermore, many (if not most) common clustering paradigms do not optimize any clearly defined objective utility, either because no such objective is defined (like in the case of, say, single linkage clustering) or because optimizing the most relevant objective is computationally infeasible, and the algorithms end up carrying out a heuristic whose outcome may be very different than the actual objective-based optimum (that is the case with the k-means algorithm as well as with spectral clustering algorithms).

We propose a different approach to providing guidance to clustering users by identifying significant properties of clustering functions that, on one hand distinguish between different clustering paradigms, and on the other hand are relevant to the domain knowledge that a user might have access to.

Based on domain expertise users could then choose which properties they want an algorithm to satisfy, and determine which algorithms satisfy each of these properties.

Our vision is that ultimately, there would be a sufficiently rich set of properties that would provide a detailed, property-based, taxonomy of clustering methods, that could, in turn, be used as guidelines for a wide variety of clustering applications. This is a very ambitious enterprise, but that should not deter researchers from addressing it. This paper takes a step towards that goal by using the natural properties to examine some popular clustering approaches.

Using our framework, we present a taxonomy for common deterministic clustering functions. We also show how to reformulate properties in the randomized setting, and use these properties to distinguish between two k-means heuristics.

Leveraging the concrete theoretical properties, we not only distinguish among various clustering algorithms, but we also study relationships between the properties, independent of any particular algorithm. In particular, we strengthen Kleinberg's impossibility result using a relaxation of one of his original properties.

1.1 Previous work

Our work follows a theoretical study of clustering that began with Kleinberg's impossibility result [8], in which he proposes three candidate axioms of clustering and shows that no clustering function can simultaneously satisfy these three axioms. Ackerman and Ben-David [1] subsequently showed these axioms to be consistent in the setting of clustering quality measures. [1] also proposes to make a distinction between clustering "axioms" and clustering "properties", where the axioms are the features that define which partitionings are worthy of the name "clustering", and the properties vary between different clustering paradigms and may be used to construct a taxonomy of clustering algorithms. We adopt that approach here.

There are previous results that provide some property based characterizations of a specific clustering algorithm. In 1975, Jardine and Sibson [6] gave a characterization of single linkage. Last year, Bosagh Zadeh and Ben-David [3] characterize single-linkage within Kleinberg's framework of clustering functions using a special invariance property ("path distance coherence"). Very recently, Ackerman, Ben-David and Loker provided characterization of the family of linkage-based clusterings in terms of a few natural properties [2].

Some heuristics have been proposed as a means of distinguishing between the output of clustering algorithms. In particular, validity criteria, or clustering quality measures, can be used to evaluate clusterings that may have been output by different algorithms. These measures can be used to select a clustering algorithm by choosing the one that yields the highest quality clustering [10]. This choice ignores the details of the clustering algorithm, relying instead on the quality measure. Unfortunately, it does not provide formal guarantees.

2 Notation and Formal Framework

Clustering is wide and heterogenous domain. For most of this paper, we focus on a basic subdomain where the (only) input to the clustering function is a finite set of points endowed with a between-points distance (or similarity) function, and the output is a partition of that domain.

A *distance function* is a symmetric function $d : X \times X \to R^+$, such that $d(x, x) = 0$ for all $x \in X$.

The data sets that we consider are pairs (X, d), where X is some finite domain set and d is a distance function over X. These are the inputs for clustering functions.

We say that a distance function d over X *extends* distance function d' over $X' \subseteq X$ if $d' \subseteq d$, that is, for all $x, y \in X'$, $d'(x, y) = d(x, y)$.

A *k-clustering* $C = \{C_1, C_2, \ldots, C_k\}$ of a data set X is a partition of X into k disjoint subsets of X (so, $\bigcup_i C_i = X$). A *clustering* of X is a k-clustering of X for some $1 \leq k \leq |X|$.

For a clustering C, let $|C|$ denote the number of clusters in C and $|C_i|$ denote the number of points in a cluster C_i. For $x, y \in X$ and a clustering C of X, we write $x \sim_C y$ if x and y belong to the same cluster in C and $x \not\sim_C y$, otherwise.

We say that (X, d) and (X', d') are *isomorphic data sets*, denoting it by $(X, d) \sim (X', d')$, if there exists a bijection $\phi : X \to X'$ so that $d(x, y) = d'(\phi(x), \phi(y))$ for all $x, y \in X$.

We consider two definitions of a clustering function. The first definition is the more general definition, a function that takes a data set and outputs a partition.

Definition 1 (General clustering function)**.** *A* general clustering function *is a function that takes as input a pair (X, d) and outputs a clustering of the domain X.*

The second type are clustering functions that require that the number of clusters be provided as part of the input. For brevity, we refer to this type of clustering functions simply as "clustering functions".

Definition 2 (Clustering function)**.** *A* clustering function *is a function that takes as input a pair (X, d) and a parameter $1 \leq k \leq |X|$ and outputs a k-clustering of the domain X.*

3 A Property-Based Characterization of Common Clustering Functions

In this section we present a taxonomy of common clustering functions. The taxonomy is presented in Figure 1 (definitions of the clustering functions are in Appendix C in the supplementary material).

A key component in our approach are properties of clustering functions that address the input-output behavior of these functions. Many such properties have recently been proposed in the literature. Kleinberg's [8] defines richness, scale-invariance, and consistency, which he proposes as axioms of clustering. We also refer to properties by Ackerman, Ben-David, and Loker [2], used in their characterization of linkage-based algorithms. Additionally, we refer to Order Invariance by Jardine and Sibson [6]. The previously-defined properties are included with the supplementary material in Appendix B, and summarized in Figure 3 in Appendix A at the end of the paper. We also propose natural variations of previous properties that allow for a more thorough study of clustering functions. The new properties are defined in Appendix A.

Except for locality and hierarchical clustering, all the properties that we discuss can be directly formulated for both clustering functions and general clustering functions[1].

The taxonomy in Figure 1 illustrates how clustering algorithms differ from one another. For example, order-invariance and inner-consistency can be used to distinguish among the three common linkage-based algorithms. Min-sum differs from K-means and K-median in that it satisfies inner-consistency. Unlike all the other algorithms discussed, the spectral clustering functions are not local.

The proofs for the taxonomy have been omitted due to a lack of space (they are included in the extra material).

3.1 Axioms of clustering

Our taxonomy reveals that some intuitive properties, that may be expected of all clustering functions, are not satisfied by some common clustering functions. For example, locality is not satisfied by the

[1]Locality can also be reformulated for general clustering functions, however, we do not focus on this in this work.

Function	outer consistent	inner consistent	local	hierarchical	order invariant	Exemplar-based	rich	outer rich	inner rich	threshold rich	scale invariant	iso. invariant
Single Linkage	✓	✓	✓	✓	✓	*X*	✓	✓	✓	✓	✓	✓
Average Linkage	✓	*X*	✓	✓	*X*	*X*	✓	✓	✓	✓	✓	✓
Complete Linkage	✓	*X*	✓	✓	✓	*X*	✓	✓	✓	✓	✓	✓
k-median	✓	*X*	✓	*X*	*X*	✓	✓	✓	✓	✓	✓	✓
k-means	✓	*X*	✓	*X*	*X*	✓	✓	✓	✓	✓	✓	✓
Min sum	✓	✓	✓	*X*	*X*	*X*	✓	✓	✓	✓	✓	✓
Ratio cut	*X*	✓	*X*	*X*	*X*	*X*	✓	✓	✓	✓	✓	✓
Normalized cut	*X*	*X*	*X*	*X*	*X*	*X*	✓	✓	✓	✓	✓	✓

Figure 1: A taxonomy of clustering functions, illustrating what properties are satisfied by some common clustering functions. The results in the k-means row apply for both versions discussed, namely, when the centers are part of the data set and when the underlying space is Euclidean and the centers are arbitrary points in the space.

spectral clustering functions ratio-cut and normalized-cut. Also, most functions fail inner consistency, and therefore do not satisfy consistency, even though the latter was previously proposed as an axiom of clustering functions [8].

On the other hand, isomorphism invariance, scale invariance, and all richness properties (in the setting where the number of clusters, k, is part of the input), are satisfied by all the clustering functions considered. Isomorphism invariance and scale-invariance make for natural axioms. Threshold richness is the only one that is both satisfied by all clustering functions considered and reflects the main objective of clustering: to group points that are close together and to separate points that are far apart.

It is easy to see that threshold richness implies richness. It can be shown that when threshold richness is combined with scale invariance, it also implies outer-richness and inner-richness. Therefore, we propose that scale-invariance, isomorphism-invariance, and threshold richness can be used as clustering axioms.

However, we emphasize that these three axioms do not make a complete set of axioms for clustering, since some functions that satisfy all three properties do not make reasonable clustering functions; a function that satisfies the two invariance properties can also satisfy threshold richness by behaving reasonably only on particularly well-clusterable data, while having counter-intuitive behavior on other data sets.

4 Properties for Randomized Clustering Functions

We present a formal setting to study and analyze probabilistic clustering functions. A *probabilistic clustering function* F takes a data set (X, d) and an integer $1 \leq k \leq |X|$ and outputs $F(X, d, k)$, a probability distribution over k-clusterings of X. Let $P(F(X, d, k) = C)$ denote the probability of clustering C in the probability distribution $F(X, d, k)$.

4.1 Properties of Probabilistic Clustering Functions

We translate properties of different types into the probabilistic setting.

Invariance properties: Invariance properties specify when data sets should be clustered in the same way (ex. isomorphism-invariance, scale-invariance, and order-invariance). Such properties are translated into the probabilistic setting by requiring that when data sets (X, d) and (X', d') satisfy some similarity requirements, then $F(X, d, k) = F(X', d', k)$ for all k.

	Properties			Axioms			Other	
Clustering Algorithm	outer consistent	local	exemplar-based	threshold rich	scale invariant	iso. invariant	rich	outer rich
Optimal k-means	✓	✓	✓	✓	✓	✓	✓	✓
Random Centroids Lloyd	X	X	✓	X	✓	✓	✓	X
Furthest Centroids Lloyd	X	X	✓	✓	✓	✓	✓	✓

Figure 2: An analysis of the k-means clustering function and k-means heuristics. The three leftmost properties are the properties that distinguish the k-means clustering function, properties that are satisfied by k-means but fail for other reasonable clustering functions. The next three are proposed axioms of clustering, and the last two properties follow from the axioms.

Consistency properties: Consistency properties impose conditions that should improve the quality of a clustering. Every such property has some notion of a "(C, d)-nice" variant that specifies how the underlying distance function can be modified to better flesh out clustering C. In the probabilistic setting, such properties require that whenever d' is a (C, d)-nice variant, the clustering function is at least as likely to output C on d' as on d, $P[F(X, d', |C|) = C] \geq P[F(X, d, |C|) = C]$.

Range properties: We consider two types of range properties, richness and range confinement properties, Richness properties require that any desired clustering can be obtained under certain constraints. In the probabilistic setting, we require that the same occurs with arbitrarily high probability. For example, the following is the probabilistic version of the richness property. The other variants of richness are reformulated analogously.

Definition 3 (Richness)**.** *A probabilistic clustering function F is* rich *if for any k-clustering C of X and any $\epsilon > 0$, there exists a distance function d over X so that $P(F(X, d, k) = C) \geq 1 - \epsilon$.*

Range confinement properties, such as exemplar-based clustering, are translated into the probabilistic setting by requiring that only those clusterings that fall within the specified range may be assigned non-zero probability by the clustering functions. We do not yet have a way to translate hierarchical clustering into the probabilistic setting.

Locality: We now show how to translate locality into the probabilistic setting. We say that a clustering of X specifies how to cluster a subset $X' \subseteq X$ if every cluster that overlaps with X' is contained within X'. Locality requires that a clustering function cluster X' in the way specified by the superset X.

In the probabilistic setting, we require that the probability of obtaining a specific clustering of $X' \subseteq X$ is determined by the probability of obtaining that clustering as a subset of $F(X, d, k)$, given that the output of F on (X, d) specifies how to cluster X'.

Definition 4 (Locality (probabilistic))**.** *A probabilistic clustering function F is* local *if for any k-clustering C' of X', $X' \subseteq X$, and $j \geq k$, where $P[\exists C_1, \ldots, C_k \text{ s.t. } \cup_{i=1}^{k} C_i = X' \mid F(X, d, j) = C] \neq 0$,*

$$P[F(X', d/X', |C'|) = C'] =$$

$$\frac{P[C' \subseteq C \mid F(X, d, j) = C \text{ and } C/X' \text{ is a } k\text{-clustering}]}{P[\exists C_1, \ldots, C_k \text{ s.t. } \cup_{i=1}^{k} C_i = X' \mid F(X, d, j) = C]}.$$

5 A Characterization of K-means Algorithms

5.1 k-means and k-means heuristics

Perhaps the most popular clustering algorithms aim to find clusterings with low k-means loss. The most popular of these use the Lloyd method. Indeed, the Lloyd method is sometimes referred to as the "k-means algorithm". We maintain a distinction between the k-means objective function and heuristics, such as the Lloyd method, which aim to find clusterings with low k-means loss. For this section, we assume that the data lie in Euclidean space, as is often the case when k-means is applied.

Definition 5 (Lloyd method)**.** *Given a data set (X, d), and a set S of points in R^n, the Lloyd algorithm performs the following steps until two consecutive iterations return the same clustering.*

1. *Assign each point in X to it's closest element of S. That is, find the clustering C of X so that $x \sim_C y$ if and only if $argmin_{c \in S} \|c - x\| = argmin_{c \in S} \|c - y\|$.*
2. *Compute the centers of mass of the clusters. Set $S = \{c_i = \frac{1}{|C_i|} \sum_{x \in C_i} x \mid C_i \in C\}$.*

The Lloyd method is highly sensitive to the choice of initial centers. Perhaps the most common method for initializing the centers for the Lloyd method is to select k random points from the input data set, proposed by Forgy in 1965 [4]. We refer to this initialization method as Random Centroids.

Katsavounidis, Kuo, and Zhang [7] proposed a deterministic method that selects centers that are far apart. Their procedure is as follows. First, let the c_1 be the point with maximum ℓ_2 norm. Then, for all $2 \leq k$, let c_i be the point furthest away from its closest existing center. That is, let c_i be the point in X that maximizes $\min_{1 \leq j \leq i-1} d(c_j, c_i)$.

We propose a variation of this initialization method, which does not depend on the embedding into Euclidean space, by setting c_1 and c_2 be the two points furthest away from each other and assigning the remaining centers as before. This is important for consistency properties, which are embedding independent. Our variation is more general, as there always exists an embedding of a data set into Euclidean space that will cause the two variations to behave the same way. The generality comes at increased computational cost of quadratic instead of linear running time in terms of the number of elements in the data set. In the remainder of the paper, we use our variation of the furthest centroids method.

5.2 Taxonomy

An analysis of the k-means clustering functions and the two k-means heuristics discussed above is shown in Figure 2. The analysis illustrates that the k-means function differs significantly from heuristics that aim to find clusterings with low k-means objective loss. The proofs for this analysis were omitted due to a lack of space (they appear in the supplementary material).

There are three properties that are satisfied by the k-means clustering function and fail for other reasonable clustering functions: outer-consistency, locality, and exemplar-based clustering. Of these, the only one satisfied by the heuristics is exemplar-based clustering, a property that is satisfied by the Lloyd method for any initialization technique.

Note that unlike clustering functions that optimize common clustering objective functions, heuristics that aim to find clusterings with low loss for these objective functions do not necessarily make meaningful clustering functions. Therefore, such heuristic's failure to satisfy certain properties does not preclude these properties from being axioms of clustering, but rather illustrates a weakness of the heuristic.

It is interesting that the Lloyd method with the Furthest Centroids initialization technique satisfies our proposed axioms of clustering while Lloyd with Random Centroid fails threshold richness. This corresponds to the finding of He et. al. [5] that in practice, Furthest Centroids performs better than Randomized Centroids.

6 Impossibility Results

In this final section, we strengthen Kleinberg's famous impossibility result [8], for general clustering functions, yielding a simpler proof of the original result.

Kleinberg impossibility theorem (Theorem 2.1, [8]) was that no general clustering function can simultaneously satisfy scale-invariance, richness, and consistency. Ackerman and Ben-David[1] later showed that consistency has some counter intuitive consequence. In Section 1, we showed that many natural clustering functions fail inner-consistency[2], which implies that there are many general clustering functions that fail consistency.

On the other hand, many natural algorithms satisfy outer-consistency. We strengthen Kleinberg's impossibility result by relaxing consistency to outer-consistency.

Theorem 1. *No general clustering function can simultaneously satisfy outer-consistency, scale-invariance, and richness.*

Proof. Let F be any general clustering function that satisfies outer-consistency, scale-invariance and richness.

Let X be some domain set with three or more elements. By richness, there exist distance functions d_1 and d_2 such that $F(X, d_1) = \{X\}$ (every domain point is a cluster on its own) and $F(X, d_2)$ is some different clustering, $C = \{C_1, \ldots C_k\}$ of X.

Let $r = \max\{d_1(x, y) : x, y \in X\}$ and let c be such that for every $x \neq y$, $cd_2(x, y) \geq r$. Define $\hat{d}(x, y) = c \cdot d_2(x, y)$, for every $x, y \in X$. Note that $\hat{d}(x, y) \geq d_1(x, y)$ for all $x, y \in X$. By outer-consistency, $F(X, \hat{d}) = F(X, d_1)$. However, by scale-invariance $F(X, \hat{d}) = F(X, d_2)$. This is a contradiction since $F(X, d_1)$ and $F(X, d_2)$ are different clusterings. □

A similar result can be obtained, using a similar proof, with inner-consistency replacing outer consistency. Namely,

Lemma 1. *No general clustering function can simultaneously satisfy inner-consistency, scale-invariance, and richness.*

Since consistency implies both outer-consistency and inner-consistency, Kleinberg's original result follows from any one of Theorem 1 or Lemma 1.

Kleinberg's impossibility result illustrates property trade-offs for general clustering functions. The good news is that these results do not apply when the number of clusters is part of the input, as is illustrated in our taxonomy; single linkage satisfies scale-invariance, consistency and richness.

7 Appendix A: Properties of Clustering Functions

The properties in this Appendix are formulated for clustering functions where the number of clusters is given as part of the input. However, all the properties, with the exception of locality and hierarchial clustering, apply for general clustering functions (in which the input is only a data set (X, d)), by a notational change from $F(X, d, k)$ to $F(X, d)$. Table 3 summarized previously defined properties. In the remainder of this appendix is used to define novel properties used in this work. These properties consist of natural variations of previous properties, and allow for a more thorough study of clustering functions.

Threshold-richness: Fundamentally, the goal of clustering is to group points that are close to each other, and to separate points that are far apart. Axioms of clustering need to represent these objectives and no set of axioms of clustering can be complete without integrating such requirements. Consistency is the only previous property that aims to formalize these requirements. However, consistency has some counterintuitive implications (see Section 3 in [1]), and is not satisfied by many common clustering functions.

Definition 6 (Threshold-richness)**.** *A clustering function F is* threshold-rich *if for every clustering C of X, there exist real numbers $a < b$ so that for every distance function d over X where*

- *$d(x, y) \leq a$ for all $x \sim_C y$, and*
- *$d(x, y) \geq b$ for all $x \not\sim_C y$*

$F(X, d, |C|) = C$.

[2]Note that a clustering function and it's corresponding general clustering function satisfy the same set of consistency properties.

Property	*Description*
Locality [2]	The behavior of the function on a union of a set of clusters is independent of the rest of the domain set
Richness [8]	Any partition of the domain can be obtained by modifying the distances between elements
Outer Richness [2]	By setting distances between some data sets, we can get the function to output each of them as a cluster (aka "extended richness")
Consistency [8]	The output remains unchanged after shrinking within-cluster distances and stretching between-cluster distances
Outer consistency [2]	The output remains unchanged after stretching between-cluster distances
Inner consistency [2]	The output remains unchanged after shrinking within-cluster distances
Scale invariance [8]	The output is invariant to uniform scaling of the data
Isomorphism [2] invariance	The output is independent of the labels of the data points (aka "representation independence")
Order invariance [6]	The output is based on the ordering of pairwise distances
Hierarchical [2]	The outputs for the same data over varying numbers of clusters are refinements of each other

Figure 3: Properties of clustering functions from the literature.

This property is based on Kleinberg's [8] Γ-forcing property, and is equivalent to the requirement that for every partition Γ, there exists $a < b$ so that (a, b) is Γ-forcing.

Inner richness: Complementary to outer richness, inner richness requires that there be a way of setting distances within sets, without modifying distances between the sets, so that F outputs each set as a cluster. This corresponds to the intuition that between-cluster distances cannot eliminate any partition of X.

Definition 7 (Inner richness)**.** *A clustering function F satisfies* inner richness *if for every data set (X, d) and partition $\{X_1, X_2, \ldots, X_k\}$ of X, there exists a $\hat{d}$ where for all $a \in X_i$, $b \in X_j$ for $i \neq j$, $\hat{d}(a, b) = d(a, b)$, and $F(\bigcup_{i=1}^{k} X_i, \hat{d}, k) = \{X_1, X_2, \ldots, X_k\}$.*

Exemplar-based clustering: Given a clustering one may be interested in finding representatives, or exemplars, for every cluster. Our new property describes clusterings whose range consists of clusterings with natural cluster exemplars.

Definition 8 (Exemplar-based)**.** *A clustering function F is* exemplar-based *if for all data sets (X, d) and $1 \leq k \leq |X|$, and every cluster $C_i \in F(X, d, k)$, there exists an exemplar e_i so that for all $x \in X$, $x \sim_C e_i$ if and only if $d(x, e_i) = \min_j d(x, e_j)$.*

Exemplars may be elements of the underlying space and lie outside of X whenever X is part of a larger domain where distances are defined (for example, if the data lies in Euclidean space). Otherwise, the exemplars are elements of the domain X.

References

[1] M. Ackerman and S. Ben-David. Measures of Clustering Quality: A Working Set of Axioms for Clustering. NIPS, 2008.

[2] M. Ackerman, S. Ben-David, and D. Loker. Characterization of Linkage-based Clustering. COLT, 2010.

[3] R. Bosagh Zadeh and S. Ben-David. "A Uniqueness Theorem for Clustering." The 25th Annual Conference on Uncertainty in Artificial Intelligence UAI, 2009.

[4] E. Forgy. Cluster analysis of multivariate data: efficiency vs. interpretability of classifications. In WNAR meetings, Univ of Calif Riverside, number 768, 1965.

[5] He, J., Lan, M., Tan, C.-L., Sung, S. -Y., and Low, H.-B. (2004). Initialization of cluster refinement algorithms: A review and comparative study. In Proc. IEEE Int. Joint Conf. Neural Networks (pp. 297-302).

[6] N. Jardine, R. Sibson, Mathematical Taxonomy Wiley, 1971.

[7] I. Katsavounidis, C.-C. J. Kuo, and Z. Zhang. A new initialization technique for generalized Lloyd iteration. IEEE Signal Processing Letters, 1(10):144146, 1994.

[8] Jon Kleinberg. "An Impossibility Theorem for Clustering." Advances in Neural Information Processing Systems (NIPS) 15, 2002.

[9] U. von Luxburg. A Tutorial on Spectral Clustering. Statistics and Computing 17(4): 395-416, 2007

[10] L. Vendramin, R.J.G.B. Campello, and E.R. Hruschka. "On the comparison of relative clustering validity criteria." Sparks, 2009.