

Statistically Rigorous Testing of Clustering Implementations

Xin Yin

Vincenzo Musco

Iulian Neamtii

Usman Roshan

*Department of Computer Science
New Jersey Institute of Technology
Newark, NJ, USA*

{xy258, vincenzo.a.musco, ineamtii, usman}@njit.edu

Abstract—Clustering is a widely-used and well-studied AI branch, but defining clustering correctness, as well as verifying and validating clustering implementations, remains a challenge. To address this, we propose a statistically rigorous approach that couples differential clustering with statistical hypothesis testing, namely we conduct statistical hypothesis testing on the outcome (distribution) of differential clustering to reveal problematic outcomes.

We employed this approach on widely-used clustering algorithms implemented in popular ML toolkits; the toolkits were tasked with clustering datasets from the Penn Machine Learning Benchmark. The results indicate that there are statistically significant differences in clustering outcomes in a variety of scenarios where users might not expect clustering outcome variation.

Index Terms—Clustering, Machine Learning, Testing, Statistics

I. INTRODUCTION

Clustering – an AI branch concerned with partitioning sets into subsets whose elements are related – is a well-established area with research going back to the 1950s. However, there is a stringent and urgent need for approaches to testing Clustering implementations due to several converging factors.

First, supervised and unsupervised learning have started to permeate software products, from “smart” home devices [1] to self-driving platforms [2] and predictive analytics [3]. These implementations make critical decisions themselves or are used to aid decision-making (e.g., autonomous driving or financial fraud detection).

Second, there has been a proliferation of clustering implementations, mostly in the form of software toolkits (e.g., MATLAB and R each offer more than 100 clustering packages [4], [5]). These implementations are run by millions of users [6], [7] including non ML-experts (from life scientists to medical professionals) who should be able to assume that the implementations are correct.

Third, software engineers are under pressure to incorporate/adopt Machine Learning into software products and processes [8]–[10]; engineers should be able to (reasonably) assume that clustering implementations are reliable and interchangeable, i.e., for a given algorithm, its implementation is correct and has no negative impact on the clustering outcome.

However, ensuring clustering correctness, or even specifying clustering correctness, remain distant goals. Therefore, we

propose an approach that can expose substantial and systematic issues with clustering algorithms’ actual implementations, e.g., wide variations across runs for theoretically-stable, deterministic algorithms, widely different outcomes for different implementations of the same algorithm, or consistently poor performance in specific implementations.

Our approach leverages the wide availability of clustering implementations and datasets with ground truth, coupled with a statistics-driven approach to help developers (or toolkit testers) find statistically significant clustering accuracy differences.

To evaluate our approach and conduct our study, we chose 7 popular clustering algorithms, 4 nondeterministic (K-means, K-means++, Spectral Clustering, Expectation Maximization-GaussianMixture); and 3 deterministic (Hierarchical clustering-Agglomerative, Affinity Propagation, DB-SCAN).¹ For uniformity and brevity, we use the following shorthands for the algorithms: *kmeans*, *kmeans++*, *gaussian*, *spectral*, *hierarchical*, *dbscan*, *apcluster*. We chose 7 widely-used clustering toolkits: MATLAB, mlpack, R, Scikit-learn, Shogun, TensorFlow, and WEKA. For uniformity and brevity, we use the following shorthands for the algorithms: *matlab*, *mlpack*, *R*, *sklearn*, *shogun*, *tensorflow*, *weka*. Our clustering inputs are 162 datasets from the Penn Machine Learning Benchmark; the datasets are described in Section II-B. 60% of the datasets come from sciences (medicine, biology, physics) with clusters crafted by domain experts.

Our basic metric of clustering similarity and accuracy is the versatile *adjusted Rand index* (ARI) [11], described in detail in Section II-A. The ARI measures the similarity between two partitions U and V of the same underlying set S . The ARI varies between -1 and $+1$, where $ARI = +1$ indicates that U and V are identical; $ARI = 0$ is tantamount to random assignment; $ARI = -1$ corresponds to strong disagreement. We use the term *accuracy* to refer to the ARI in the case when U is a clustering produced by an algorithm and V is the Ground Truth (as labeled in PMLB) for that dataset.

Examples: clustering accuracy in promoters. The clustering accuracy is measured by ARI between clustering results and ground truth. The promoters [12] dataset essentially contains

¹Deterministic clustering algorithms should, in theory, produce the same clustering when run repeatedly on the same input. For nondeterministic (aka randomized) algorithms, the clustering might vary.




Fig. 1. Kmeans++: Accuracy distributions for 11 toolkits on dataset promoters (WN = WEKA with default normalization).

two clusters that partition the underlying E.coli DNA (gene) sequences into “promoters”² and “non-promoters”. We ran our statistical approach on clustering outcomes of algorithm K-means++ in 11 toolkit configurations (sample size = 30, corresponding to the 30 repeated runs of that configuration on the same dataset). The resulting accuracy range is shown in Figure 1. Note that 10 toolkits achieve identical, perfect clusterings in each of the 30 runs, i.e.,

$$\min Accuracy = \max Accuracy = 1$$

However, the toolkit WEKA’s accuracy across 30 runs varied:

$$\min Accuracy = -0.01; \max Accuracy = 0.4$$

This allows us to make two observations: (1) WEKA’s accuracy varies substantially across re-runs on the same input dataset; and (2) all toolkits achieve perfect clusterings in every run, whereas’s WEKA’s best run has accuracy 0.4. This clearly indicates a clustering implementation issue.³

Our approach exposes such issues (and more) in a statistically rigorous way. *This is the first statistically rigorous approach to testing clustering implementations.*

Our approach has 4 components; we present each component in the context of using that component to test a null hypothesis:

- 1) How different runs of the same algorithm in the same implementation lead to different clusterings (Section III).
- 2) How different implementations of the same algorithm in different toolkits lead to different clusterings (Section IV).
- 3) The toolkit’s impact when comparing algorithms (Section V).
- 4) How different toolkits “disagree” (Section VI).

²A *promoter* sequence marks the DNA region where the transcription of that gene into RNA should begin.

³WEKA developers were able to attribute this low accuracy to WEKA’s default normalization setting.

TABLE I
CATEGORIES FOR THE PMLB DATASETS.

Category	Percentage
Medical/Health	24%
Biology, Biochemistry, Bioinformatics	15%
Physics, Math, Astronomy	11%
Social, Census	10%
Sports	7%
Financial	7%
Image recognition	6%
Synthetic datasets	6%
IT, AI	4%
Linguistics	3%
Miscellaneous	7%

II. BACKGROUND

A. Definitions

Clustering. Given a set S of n points (d -dimensional vectors in the \mathcal{R}^d space), a clustering is a partitioning of S into K non-overlapping subsets (clusters) $S_1, \dots, S_i, \dots, S_K$ such that intra-cluster distance between points (that is, within individual S_i ’s) is minimized, while inter-cluster distance (e.g., between centroids of S_i and S_j where $i \neq j$) is maximized.

Accuracy. The *adjusted Rand index* (ARI) is a versatile, widely used [13] clustering comparison metric. Let us assume a set D and two partitionings A and B on it: P_A, Q_A (where $P_A \cup Q_A = D$ and $P_A \cap Q_A = \emptyset$) and respectively P_B, Q_B (where $P_B \cup Q_B = D$ and $P_B \cap Q_B = \emptyset$). We can measure these two clustering’s similarity using the ARI. When $ARI = +1$, we have $P_A = P_B$ and $P_B = Q_B$. When $ARI \neq +1$, the two partitioning schemes differ. The case $ARI = 0$ corresponds to D elements being assigned randomly across P_A/P_B and Q_A/Q_B , respectively. The case $-1 < ARI < 0$ is defined as “worse-than-random”, because it is worse than randomness clustering. Technically, the ARI of A and B is defined:⁴

$$ARI(A, B) = \frac{2(N_{00}N_{11} - N_{01}N_{10})}{(N_{00} + N_{01})(N_{01} + N_{11}) + (N_{00} + N_{10})(N_{10} + N_{11})}$$

B. Datasets

We chose 162 datasets from the 166-dataset PMLB (Penn Machine Learning Benchmark) [15], a benchmark suite that includes “real-world, simulated, and toy benchmark datasets” [15].⁵ PMLB was designed to benchmark ML implementations and avoid imbalance across meta-features (which often plagues handpicked datasets). The geometric means across the dataset collection – roughly, the characteristics of the typical dataset – were 809 instances, 15.41 features, and 3.18 clusters.

⁴Where “ N_{11} is the number of pairs that are in the same cluster in both A and B ; N_{00} is the number of pairs that are in different clusters in both A and B ; N_{01} is the number of pairs that are in the same cluster in A but in different clusters in B ; and N_{10} is the number of pairs that are in different clusters in A but in the same cluster in B ” [14].

⁵For efficiency, we set aside four large datasets: connect-4, poker, mnist, kddcup.

C. Clustering Algorithms

We now present a brief overview of each algorithm.

K-means. The algorithm aims to cluster the observations (points in S) into K distinct clusters, where observations belong to the clusters with the nearest mean. The goal is to minimize the sum of all intra-cluster distances. The algorithm starts from K selected initial points as “centroids” (cluster centers). These centroids play a crucial role in the algorithm’s effectiveness: the algorithm is not guaranteed to converge to a global minimum, so with “bad” centroids the algorithm can “fall” into local minima.

Variation due to starting points: the algorithm requires “starting points”, that is, initial centroids. To ensure consistency across toolkits, we fed all toolkits the same starting points. We explored the variation in outcome by randomly picking different starting points from the datasets. That is, if a dataset has $K = 2$, for each run we pick two different points in S , s_1 and s_2 as centroids, and run the algorithm from there, with all toolkits starting with s_1 and s_2 .

K-means++ was designed to improve K -means by choosing the starting points more carefully so they are farther apart. Theoretically, this improved version ensures that the algorithm is less likely to converge to local minima.

EM/Gaussian. Gaussian mixture clustering is a model-based approach: clustering is first done using a model (i.e., a parametric distribution such as a Gaussian). Each cluster is defined by an initial random model and the dataset is composed of the mixture of these models. Then, the model/data fitness is optimized – a common optimization is Expectation-Maximization.

Spectral clustering computes eigenvalues of the similarity matrix between the data points to reduce the dimensionality of the original data. After the reduction, a clustering algorithm, e.g., K -means, is applied on the reduced-dimensionality data.

The aforementioned 4 algorithms were nondeterministic; we now discuss the 3 deterministic algorithms.

Hierarchical clustering is a deterministic algorithm, based on building a hierarchy of clusters using one of two approaches: (i) a bottom-up approach named “agglomerative”: each observation is initially put in its own cluster and then they are merged, (ii) a top-down approach named “divisive”: all observations are initially placed in the same cluster and then they are split. We use the first variant.

DBSCAN is a deterministic algorithm based on density, that is, high density regions of data are grouped together in a neighbor graph and form a data cluster.

Affinity Propagation (AP) is based on propagation on a graph in which each data point is a node. The algorithm builds clusters by iteratively passing messages from a node to another until determining which one is part of a specific cluster.

Some toolkits do not support all 7 algorithms; Table II shows the supported algorithm/toolkit combinations; in all, there were 27 algorithm-toolkit configurations.

TABLE II
TOOLKIT/ALGORITHM CONFIGURATIONS

	MATLAB	mlpack	R	Scikit-learn	Shogun	TensorFlow	WEKA
kmeans++	✓	✓	✓	✓	✓	✓	✓
kmeans	✓	✓	✓	✓	✓	✓	✓
spectral			✓	✓			
hierarchical	✓		✓	✓			
gaussian	✓			✓		✓	✓
dbscan		✓	✓	✓			
apcluster			✓	✓			

TABLE III
LEVENE’S TEST RESULTS: THE NUMBER OF DATASETS, OUT OF 162, WITH SIGNIFICANT VARIANCE ($p < 0.05$)

Algorithm	Toolkit	# Datasets
kmeans++	sklearn	126
kmeans++	R	111
kmeans++	mlpack	144
kmeans++	matlab	125
kmeans++	shogun	143
kmeans++	tensorflow	144
kmeans++	weka	157
spectral	sklearn	93
spectral	sklearnfast	97
spectral	R	113
kmeans	sklearn	148
kmeans	R	153
kmeans	mlpack	146
kmeans	matlab	141
kmeans	shogun	146
kmeans	tensorflow	145
hierarchical	sklearn	71
hierarchical	R	63
hierarchical	matlab	63
gaussian	sklearn	136
gaussian	matlab	153
gaussian	tensorflow	151
gaussian	weka	123

D. Runs

Our analysis is based on clustering results achieved by each algorithm-toolkit configuration on each of the 162 datasets 30 different times (i.e., more than 160,000 clustering tasks; for all toolkits default settings were used).

Specifically we analyzed the distribution of accuracy (i.e., ARI when comparing the resulting clustering with ground truth) achieved in these 30 different clustering runs. Hence for our subsequent statistical analyses we have sample size $n = 30$ for one-sample tests (and $n_1 = n_2 = 30$ for two-sample tests). Note that K -means requires “starting points” – initial centroids. Hence in configuration *kmeans* each of the 30 runs used a different, randomly-picked, different starting points from the dataset.

III. VARIATION ACROSS RUNS

This testing procedure is shown in Figure 2: a single toolkit is run on a single dataset multiple times (30 in our case), and a statistical analysis is performed on the resulting accuracy distribution.

TABLE IV
TOP-10 WIDEST DIFFERENCES IN ACCURACY ACROSS RUNS

Algorithm	Toolkit	Dataset	Min	Max
gaussian	tensorflow	prnn_crabs	-0.005	1
gaussian	matlab	prnn_crabs	-0.005	0.979
gaussian	matlab	anacatdata_cr.	-0.024	0.958
gaussian	tensorflow	twonorm	0	0.908
gaussian	tensorflow	twonorm	0.003	0.910
gaussian	tensorflow	ionosphere	0.004	0.772
spectral	R	breast-w	0.056	0.818
gaussian	matlab	anacatdata_aut.p	0.041	0.794
gaussian	matlab	wdbc	0.007	0.754
gaussian	tensorflow	breast-cancer-wsc.	0.032	0.760

TABLE V
HIGHEST STANDARD DEVIATIONS IN ACCURACY ACROSS RUNS

Algorithm	Toolkit	Dataset	Stddev
gaussian	tensorflow	twonorm	0.400
gaussian	tensorflow	prnn_crabs	0.345
gaussian	tensorflow	ionosphere	0.298
gaussian	sklearn	breast	0.281
kmeans++	weka	australian	0.236
gaussian	matlab	house-votes-84	0.236
gaussian	matlab	tokyo1	0.216
gaussian	sklearn.0_tol	breast	0.213
gaussian	matlab	twonorm	0.212
gaussian	matlab	anacatdata_cr.	0.206
gaussian	matlab	wine-recognition	0.205
spectral	R	appendicitis	0.204
kmeans++	shogun	house-votes-84	0.201

Null hypothesis: accuracy does not vary across runs.

In other words, for a certain algorithm and dataset, we set out to measure *non-determinism*. To test this hypothesis, we use Levene’s test as follows: one sample contains the actual accuracy values for the 30 runs, the other sample has the same mean, size, and no variance, that is, all 30 elements are equal to the mean of the first set. We ran this on all datasets. Rejecting the null hypothesis means that accuracy varies in a statistically significant way across runs. We report results at $p < 0.05$.

In Table III we show the number of datasets where variance is statistically significant at $p < 0.05$; recall that we have a total of 162 datasets. We observe that Spectral is the most stable nondeterministic algorithm; for Spectral, only 93–113 datasets show significant variance. Hierarchical, which should be deterministic, still has 63–71 datasets with significant variance. In contrast, *K-means*, *K-means++*, and *Gaussian Mixture*, have significant variance from run to run.

In Table IV we show how broad the accuracy range (difference between minimum accuracy and maximum accuracy) can be. The first three columns show the algorithm, toolkit and dataset. The last two columns show the minimum and maximum accuracy attained over the 30 runs. For example, Gaussian has quite a large range on some datasets: accuracy on the dataset *prnn_crabs* has a min-max range of *more than 1*, with one run’s accuracy below 0 and another run having perfect or (close to perfect) accuracy.

In Table V we show how high the standard deviation of the accuracy can be across runs. For example, accuracy on the dataset *twonorm* can have a *standard deviation of 0.4*. More

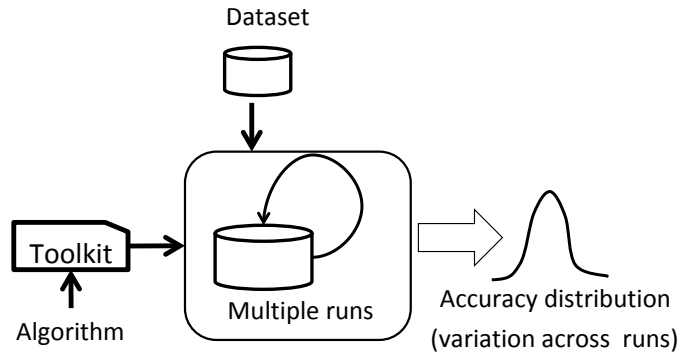


Fig. 2. Testing for variation across runs.

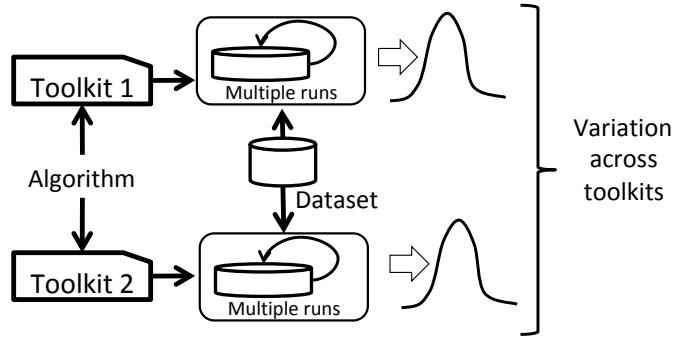


Fig. 3. Testing for variation across toolkits.

than a dozen other toolkit/algorithm setups have *standard deviation higher than 0.2*.

IV. VARIATION ACROSS TOOLKITS

This testing procedure is shown in Figure 3: two toolkits implementing the same algorithms are run on the same dataset multiple times (30 in our case), and a statistical analysis is performed to compare the two accuracy distributions.

Null hypothesis: For a given algorithm, accuracy does not vary across toolkits.

To test this hypothesis, we use the Mann-Whitney U test as follows. We fix the algorithm, e.g., *K-means++*, and the dataset. Next, we compare the distributions of accuracy values pairwise, for all possible toolkits pairs, that is, if we have N toolkits for a given algorithm, for a given dataset there will be $\binom{N}{2}$ Mann-Whitney U tests; hence for each algorithm there will be $162 \times \binom{N}{2}$ tests. Rejecting the null hypothesis means that accuracy varies significantly between toolkits. We report results at $p < 0.05$.

In Table VI we show the number of datasets where accuracy distributions between two toolkits are statistically significant at $p < 0.05$. We observe that *Gaussian Mixture* and *K-means++* induce most differences in toolkit outcomes’ distributions (generally over 100 out of 162). Even for *apcluster* (deterministic), on 40 out of 162 datasets we found statistically significant differences between *Sklearn* and *R*.

TABLE VI
MANN-WHITNEY U-TEST RESULTS FOR TOOLKITS: NUMBER OF DATASETS WITH SIGNIFICANTLY DIFFERENT ACCURACY DISTRIBUTIONS ($p < 0.05$)

Algorithm	Toolkits	# Datasets
kmeans++	sklearn vs. R	50
kmeans++	sklearn vs. matlab	107
kmeans++	sklearn vs. weka	134
kmeans++	sklearn vs. mlpack	104
kmeans++	sklearn vs. shogun	110
kmeans++	sklearn vs. tensorflow	109
kmeans++	R vs. matlab	104
kmeans++	R vs. weka	134
kmeans++	R vs. mlpack	101
kmeans++	R vs. shogun	108
kmeans++	R vs. tensorflow	115
kmeans++	matlab vs. weka	141
kmeans++	matlab vs. mlpack	105
kmeans++	matlab vs. shogun	120
kmeans++	matlab vs. tensorflow	124
kmeans++	weka vs. mlpack	96
kmeans++	weka vs. shogun	113
kmeans++	weka vs. tensorflow	125
kmeans++	mlpack vs. shogun	15
kmeans++	mlpack vs. tensorflow	57
kmeans++	shogun vs. tensorflow	60
spectral	sklearn vs. R	109
kmeans	sklearn vs. R	41
kmeans	sklearn vs. matlab	9
kmeans	sklearn vs. mlpack	8
kmeans	sklearn vs. shogun	9
kmeans	sklearn vs. tensorflow	9
kmeans	R vs. matlab	48
kmeans	R vs. mlpack	39
kmeans	R vs. shogun	43
kmeans	R vs. tensorflow	42
kmeans	matlab vs. mlpack	2
kmeans	matlab vs. shogun	1
kmeans	matlab vs. tensorflow	0
kmeans	mlpack vs. shogun	1
kmeans	mlpack vs. tensorflow	2
kmeans	shogun vs. tensorflow	1
hierarchical	sklearn vs. R	53
hierarchical	sklearn vs. matlab	58
hierarchical	R vs. matlab	57
gaussian	sklearn vs. matlab	129
gaussian	sklearn vs. weka	146
gaussian	sklearn vs. tensorflow	120
gaussian	matlab vs. weka	146
gaussian	matlab vs. tensorflow	104
gaussian	weka vs. tensorflow	120
dbscan	sklearn vs. R	0
dbscan	sklearn vs. mlpack	7
dbscan	R vs. mlpack	7
apcluster	sklearn vs. R	40

A. Non-overlaps

In Table VII we show the largest gaps between accuracy intervals, computed as follows: we find all dataset/algorithm combinations where the accuracy intervals for two toolkits, say $[ARI_{1min}, ARI_{1max}]$ and $[ARI_{2min}, ARI_{2max}]$ are non-overlapping, that is, $ARI_{1min} > ARI_{2max}$. In other words, for any run of toolkit 1, its accuracy floor (min.) is higher than the accuracy ceiling (max.) of any run of toolkit 2. We call that difference “gap”, i.e., $gap = ARI_{1min} - ARI_{2max}$. We show the top-10 gaps in Table VII. Notice that this gap can be as large as 0.966.

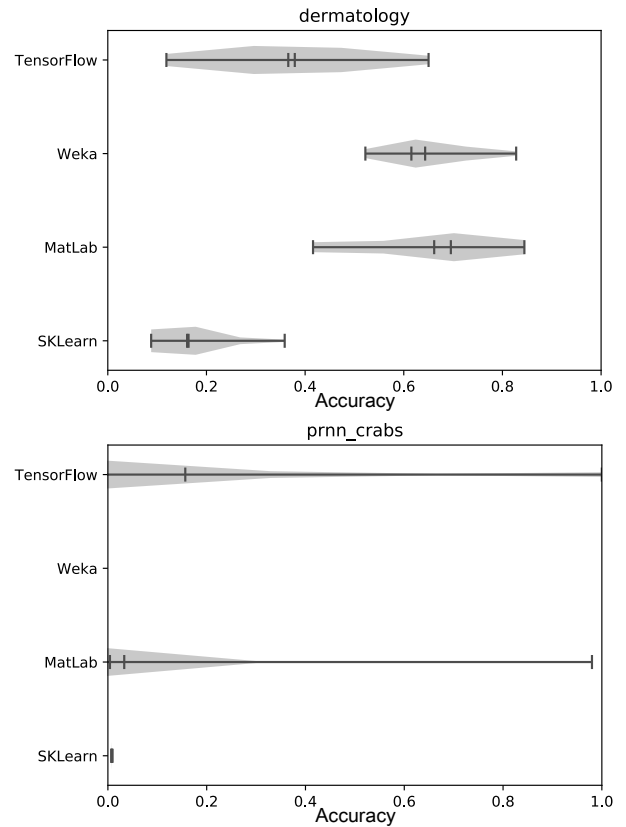


Fig. 4. EM (Gaussian Mixture): differences between toolkits on two datasets, dermatology and prnn-crabs.




Fig. 5. Testing for variation across algorithms.

We found that 1,776 such gaps exist (out of 34,987 runs of the same algorithm/dataset combinations). This is very problematic, as it shows how toolkits are not “created equal” – even after multiple runs, in 1,776 scenarios, a toolkit’s best accuracy cannot even reach another toolkit’s worst accuracy.

In Figure 4 we show violin plots of toolkits’ accuracy distributions in the EM algorithm on two datasets. On set dermatology (top) note the wide ranges of TensorFlow and the gap between WEKA and Sklearn. On set prnn-crabs (bottom) note the high-end accuracy of 1 (TensorFlow, MATLAB) and the consistently low accuracy in WEKA and Sklearn.

V. VARIATION ACROSS ALGORITHMS

This testing procedure is shown in Figure 5: implementations of two different algorithms but in the same toolkit are

TABLE VII
TOP-10 LARGEST ACCURACY GAPS BETWEEN TOOLKITS

Algorithm	Dataset	Toolkit 1		Toolkit 2		Gap
			Floor (Min)		Ceiling (Max)	
gaussian	promoters	sklearn	1	tensorflow	0.034	0.966
gaussian	promoters	sklearn	1	tensorflow	0.034	0.966
spectral	promoters	R	0.962	sklearn	0.001	0.962
spectral	analcatdata_cred.	sklearn	0.84	R	-0.002	0.842
kpp	breast	weka	0.813	sklearn	-0.003	0.815
kpp	breast	weka	0.813	mlpack,matlab,R	0.02	0.792
kpp	breast	weka	0.813	tensorflow,shogun	0.02	0.792
gaussian	promoters	sklearn	1	matlab	0.234	0.766
kpp	promoters	tensorflow	1	weka	0.406	0.594

run on the same dataset multiple times (30 in our case), and a statistical analysis is performed to compare the two accuracy distributions.

Null hypothesis: For a given toolkit, accuracy does not vary across algorithms.

To test this hypothesis, we again use the Mann-Whitney U test. We fix the toolkit, e.g., MATLAB, and the dataset. Next, we compare the distributions of accuracy values pairwise, for all possible algorithm pairs. Rejecting the null hypothesis implies that, for a given toolkit, algorithms’ accuracy varies significantly.

In Table VIII we show the number of datasets where accuracy distributions between two algorithms are significantly different. Typically, algorithms’ accuracies differ on more than 110 datasets; we expected to see such differences between algorithms. However, we did not expect wide differences when looking at the *same algorithm pairs in different toolkits*. For example, K -means and K -means++ differ on 105/101/115/120 datasets in Sklearn/R/MATLAB/WEKA but only on 27 datasets in MLpack and only 31 datasets in Shogun. This again shows that toolkits are not interchangeable (though users might expect them to be).

VI. TOOLKIT DISAGREEMENT

We next set out to study whether toolkits “agree” or “disagree” on those points that are misclassified w.r.t. ground truth. Specifically, we are interested in those cases where two toolkits have relatively high accuracy w.r.t. ground truth, but there are large disagreements between the toolkits on the remaining, or misclassified points (i.e., where toolkits’ clustering differs from ground truth).

We illustrate this in Figure 6. Assuming two toolkits $T1$ and $T2$, their clustering of x_1, x_2, x_3, x_4 (on the bottom) is in agreement, and let us assume this clustering agrees with ground truth as well. We want to measure the disagreement on the remaining points x_5, x_6, x_7, x_8 (on top).

Intuitively, datasets that induce this disagreement between $T1$ and $T2$ on the top points manage to expose differences in toolkit implementations “at the margin”; since agreement with ground truth is high, users might expect the toolkits will be in agreement on the remaining points, too.

Let ARI_{T1G} and ARI_{T2G} be the accuracy of two different toolkits on the same algorithm and same dataset. Let ARI_{T1T2}




Fig. 6. Toolkit disagreement.

be the ARI when comparing the two clusterings (rather than with ground truth). There were 14,831 ARI_{T1T2} comparisons. Out of these, we found 928 cases where:

$$ARI_{T1G} > ARI_{T1T2} \wedge ARI_{T2G} > ARI_{T1T2}$$

That is, there were 928 cases where toolkits’ clusterings disagree with each other more than they disagree with ground truth – in other words, toolkits disagree strongly on those points that are not clustered perfectly.

In Table IX we show the top-10 disagreements, excluding the trivial cases where one toolkit’s accuracy is 1. These datasets are particularly important as they manage to “drive wedges” between toolkits; this has many applications, from differential toolkit testing to constructing adversarial datasets.

VII. RELATED WORK

There is a surprising scarcity of studies measuring clustering accuracy/outcome across toolkits and across same-toolkit runs.

Fränti [16] has compared performance on clustering basic benchmark, and measure performance on four factors: overlap of clusters, number of clusters, dimensionality and unbalance of cluster sizes. However, they only consider synthesis data and their datasets have simple structures. Our work is based on PMLB which includes mainly real-world datasets allowed for comparing ML methods comprehensively.

Hamerly [17] has proposed a new algorithm for accelerating K -means, and performed an evaluation on efficiency similar to

TABLE VIII
MANN-WHITNEY U-TEST RESULTS FOR ALGORITHMS: NUMBER OF DATASETS WITH SIGNIFICANTLY DIFFERENT ACCURACY DISTRIBUTIONS ($p < 0.05$)

Toolkit	Algorithms	# Datasets
sklearn	kmeans vs. kmeans++	105
sklearn	kmeans vs. gaussian	123
sklearn	kmeans vs. hierarchical	134
sklearn	kmeans vs. spectral	112
sklearn	kmeans vs. dbscan	150
sklearn	kmeans vs. apcluster	115
sklearn	kmeans++ vs. gaussian	132
sklearn	kmeans++ vs. hierarchical	153
sklearn	kmeans++ vs. spectral	109
sklearn	kmeans++ vs. dbscan	155
sklearn	kmeans++ vs. apcluster	117
sklearn	gaussian vs. hierarchical	145
sklearn	gaussian vs. spectral	108
sklearn	gaussian vs. dbscan	150
sklearn	gaussian vs. apcluster	113
sklearn	hierarchical vs. spectral	120
sklearn	hierarchical vs. dbscan	155
sklearn	hierarchical vs. apcluster	122
sklearn	spectral vs. dbscan	122
sklearn	spectral vs. apcluster	115
sklearn	dbscan vs. apcluster	117
shogun	kmeans vs. kmeans++	31
R	kmeans vs. kmeans++	101
R	kmeans vs. hierarchical	139
R	kmeans vs. spectral	94
R	kmeans vs. dbscan	149
R	kmeans vs. apcluster	117
R	kmeans++ vs. hierarchical	150
R	kmeans++ vs. spectral	97
R	kmeans++ vs. dbscan	157
R	kmeans++ vs. apcluster	123
R	hierarchical vs. spectral	113
R	hierarchical vs. dbscan	154
R	hierarchical vs. apcluster	123
R	spectral vs. dbscan	115
R	spectral vs. apcluster	122
R	dbscan vs. apcluster	123
tensorflow	kmeans vs. kmeans++	74
tensorflow	kmeans vs. gaussian	117
tensorflow	kmeans++ vs. gaussian	121
matlab	kmeans vs. kmeans++	115
matlab	kmeans vs. gaussian	135
matlab	kmeans vs. hierarchical	141
matlab	kmeans++ vs. gaussian	146
matlab	kmeans++ vs. hierarchical	155
matlab	gaussian vs. hierarchical	146
mlpack	kmeans vs. kmeans++	27
mlpack	kmeans vs. dbscan	135
mlpack	kmeans++ vs. dbscan	130
weka	kmeans++ vs. gaussian	120

Kriegel et al.’s (time and memory). Our focus is on accuracy rather than efficiency.

Chen et al. [18] have compared four clustering algorithms – hierarchical clustering, K -means, Self-organizing Map (SOM) and Partitioning around Medoids (PAM) on a single dataset, mouse genomic data. Unlike us, they varied the K , whereas we used the ground truth’s K . Our focus is different: varying runs of the same algorithm, and a breadth of datasets. Abu [19] has compared four clustering algorithms – K -means, hierarchical, SOM, and Expectation Maximization (EM), each implemented in two toolkits LNKnet and Cluster/TreeView; they used a sin-

gle 600-instance dataset, and compared performance/accuracy on this dataset, and a 200-instance subset thereof. Our setup is substantially larger and our focus is substantially broader.

Clustering stability has been defined by Tilman et al. [20] as “solutions [that] are similar for two different data sets that have been generated by the same (probabilistic) source”. Our definition of stability is different: similarity of solutions on the same dataset, but produced by two different runs.

Our own work on SmokeOut [21] also explored the space of different clustering outcomes across toolkits and runs, but the goal was to expose outliers and characterize distributions shapes. Descriptive statistics (min/max) were used for the comparison, which is concise but lacks statistical rigor. In contrast, in this paper we introduce and use a statically rigorous approach for comparing runs, toolkits and algorithms. SmokeOut was also run on PMLB. Another goal of SmokeOut was to expose outliers or bimodal distributions in clustering outcomes across repeated runs – information which is useful for debugging.

VIII. CONCLUSIONS

We present the first approach for testing clustering implementations via rigorous statistical analysis. We demonstrate our approach via statistical analysis of clustering outcomes across multiple runs, toolkits and algorithms. We found statistically significant variations across all these dimensions, which might violate users’ determinism and invariance assumptions. Our results point out the need for improving the correctness and determinism of clustering implementations.

ACKNOWLEDGMENTS

This material is based upon work supported by the NSF Grant No. CCF-1629186. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] D. Bell, “5 great ai-powered home devices that will improve your life today,” <https://www.t3.com/features/5-great-ai-powered-home-devices-that-will-improve-your-life-today>.
- [2] Nvidia, “World’s first functionally safe ai self-driving platform,” <https://www.nvidia.com/en-us/self-driving-cars/drive-platform/>.
- [3] PAT RESEARCH, “Top 15 artificial intelligence platforms in 2018,” 2018, <https://www.predictiveanalyticstoday.com/artificial-intelligence-platforms/>.
- [4] “Cran task view: Cluster analysis & finite mixture models,” November 2018, <https://cran.r-project.org/web/views/Cluster.html>.
- [5] “Matlab file exchange:clustering,” November 2018, <https://www.mathworks.com/matlabcentral/fileexchange/?term=clustering+product%3A%22MATLAB%22&utf8=%E2%9C%93>.
- [6] “Mathworks fast facts,” <https://www.mathworks.com/company/aboutus.html>.

TABLE IX
TOP-10 LARGEST DISAGREEMENTS BETWEEN TOOLKITS YET HAVING HIGH AGREEMENT WITH GROUND TRUTH

Algorithm	Dataset	Toolkit1		Toolkit2		
			ARI_{T1G}		ARI_{T2G}	ARI_{T1T2}
spectral	promoters	sklearnfast	0.889	R	0.962	0.853
gaussian	iris	weka	0.759	sklearn	0.904	0.693
gaussian	wine-recognition	weka	0.915	sklearn	0.607	0.568
gaussian	analcatdata_authorship	weka	0.951	sklearn	0.740	0.719
gaussian	wine-recognition	sklearn	0.607	matlab	0.724	0.469
spectral	breast-w	sklearnfast	0.809	R	0.552	0.477
gaussian	wine-recognition	weka	0.915	matlab	0.724	0.718
gaussian	iris	sklearn	0.904	matlab	0.560	0.550
gaussian	iris	tensorflow	0.562	sklearn	0.904	0.555
gaussian	texture	tensorflow	0.694	sklearn	0.742	0.614
gaussian	dermatology	weka	0.615	matlab	0.695	0.519
kpp	analcatdata_authorship	weka	0.777	shogun	0.718	0.700

- [7] M. Hornick, "Oracle r technologies overview," <https://www.oracle.com/assets/media/oraclertechnologies-2188877.pdf>.
- [8] Janakiram MSV, "Why do developers find it hard to learn machine learning?" 2017, <https://www.forbes.com/sites/janakirammsv/2018/01/01/why-do-developers-find-it-hard-to-learn-machine-learning/>.
- [9] Carlton E. Sapp, "Gartner: Preparing and architecting for machine learning," 2017, https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_and_architecting_for_machine_learning.pdf.
- [10] N. M. do Nascimento, C. Lucena, P. S. C. Alencar, and D. D. Cowan, "Software engineers vs. machine learning algorithms: An empirical study assessing performance and reuse tasks," *CoRR*, vol. abs/1802.01096, 2018.
- [11] L. Hubert and P. Arabie, "Comparing partitions," vol. 2, pp. 193–218, 02 1985.
- [12] D. Dheeru and E. Karra Taniskidou, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [13] D. Steinley, "Properties of the hubert-arable adjusted rand index." *Psychological methods*, vol. 9, no. 3, p. 386, 2004.
- [14] N. X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance," *JMLR*, vol. 11, no. Oct, pp. 2837–2854, 2010.
- [15] R. S. Olson, W. La Cava, P. Orzechowski, R. J. Urbanowicz, and J. H. Moore, "Pmlb: a large benchmark suite for machine learning evaluation and comparison," *BioData Mining*, vol. 10, no. 1, p. 36, Dec 2017.
- [16] P. Fránti and S. Sieranoja, "K-means properties on six clustering benchmark datasets," *Applied Intelligence*, vol. 48, no. 12, pp. 4743–4759, Dec 2018. [Online]. Available: <https://doi.org/10.1007/s10489-018-1238-7>
- [17] G. Hamerly, *Making k-means even faster*, pp. 130–140. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/1.9781611972801.12>
- [18] G. Chen, S. A. Jaradat, N. Banerjee, T. S. Tanaka, M. S. H. Ko, and M. Q. Zhang, "Evaluation and comparison of clustering algorithms in analyzing es cell gene expression data," *Stat. Sinica*, pp. 241–262, 2002.
- [19] O. Abu Abbas, "Comparison between data clustering algorithm," *Int. Arab Journal of Information Technology*, vol. 5, no. 3, 2008.
- [20] T. Lange, V. Roth, M. L. Braun, and J. M. Buhmann, "Stability-based validation of clustering solutions," *Neural Computation*, vol. 16, no. 6, pp. 1299–1323, 2004.
- [21] V. Musco, X. Yin, and I. Neamtiu, "Smokeout: An approach for testing clustering implementations," in *12th IEEE International Conference on Software Testing, Verification and Validation, ICST 2019*, 2019, to appear.