

Accuracy Evaluation of Overlapping and Multi-resolution Clustering Algorithms on Large Datasets

Artem Lutov, Mourad Khayati and Philippe Cudré-Mauroux
eXascale Infolab, University of Fribourg—Switzerland
Email: {firstname.lastname}@unifr.ch

Abstract—Performance of clustering algorithms is evaluated with the help of accuracy metrics. There is a great diversity of clustering algorithms, which are key components of many data analysis and exploration systems. However, there exist only few metrics for the accuracy measurement of overlapping and multi-resolution clustering algorithms on large datasets. In this paper, we first discuss existing metrics, how they satisfy a set of formal constraints, and how they can be applied to specific cases. Then, we propose several optimizations and extensions of these metrics. More specifically, we introduce a new indexing technique to reduce both the runtime and the memory complexity of the Mean F1 score evaluation. Our technique can be applied on large datasets and it is faster on a single CPU than state-of-the-art implementations running on high-performance servers. In addition, we propose several extensions of the discussed metrics to improve their effectiveness and satisfaction to formal constraints without affecting their efficiency. All the metrics discussed in this paper are implemented in C++ and are available for free as open-source packages that can be used either as stand-alone tools or as part of a benchmarking system to compare various clustering algorithms.

Index Terms—accuracy metrics, overlapping community evaluation, multi-resolution clustering evaluation, Generalized NMI, Omega Index, MF1, similarity of collections of sets

I. INTRODUCTION

Clustering is a key component of many data mining systems with numerous applications including statistical analysis and the exploration of physical, social, biological and informational systems. This diversity of potential applications spawned a wide variety of network (graph) clustering algorithms proposed in the literature. It also led to specialized clustering algorithms in particular domains. Hence, the need to find the most suitable and best performing clustering algorithms for a given task became more dire, and the evaluation of the resulting clustering through proper metrics more important. Moreover, as modern systems often operate on very large datasets (consisting of billions of items potentially), the computational properties of the evaluation metrics become more important. In particular, performance-related constraints rapidly emerge when sampling the original large datasets is not desirable or not possible for a given use-case.

Clustering quality metrics can formally be categorized into two types: intrinsic and extrinsic metrics. Intrinsic quality metrics evaluate how the elements of each cluster are similar to each other and how they differ from elements in other clusters given a similarity metric such as modularity [1] or conductance [2]. Extrinsic quality metrics (also known as *accuracy* metrics) evaluate instead how the clusters are similar to the *ground-truth* (gold standard) clusters. In this paper, we focus on extrinsic metrics since they allow to identify clustering algorithms producing expected results, which is in practice often more useful than measuring the formation of optimal clusters by a given similarity metric. More specifically, we evaluate the similarity (proximity) between two collections of overlapping clusters (unordered sets) of elements, where *a*) the collections have the same number of elements, *b*) each element may be a member of multiple clusters and *c*) each cluster may have several, non-mutual, best (in terms of the similarity value) matches in the other collection.

We call *clustering* the set of clusters resulting from an algorithm. Clusterings can be categorized as non-overlapping (crisp clustering, hard partitioning), overlapping (soft, fuzzy clustering) or, in some cases, multi-resolution (including hierarchical). Multi-resolution clusterings are considered when there is a need to simultaneously compare multiple resolutions (hierarchy levels) of the results against the ground-ground, where each resolution contains non/overlapping clusters as discussed further in Section IV-D. Non-overlapping clusterings can be seen as a special case of overlapping clusterings.

A large number of accuracy metrics were proposed in the literature to measure the clustering quality [3]–[8]. Evaluation frameworks [9]–[13] and surveys [14]–[17] were also introduced. Despite the large number of accuracy metrics proposed, very few metrics are applicable to overlapping clusters, causing many issues when evaluating such clusters (e.g., Adjusted Rand Index is used to evaluate overlapping clusters in [18], even though it is applicable to non-overlapping clusters only). Moreover, most of the quality metrics for overlapping clusters are not comparable to similar metrics for non-overlapping clusters (e.g., standard NMI [5] or modularity [1] versus some overlapping NMI [19] or overlapping modularity [20]–[23] implementations), which complicates the direct comparison of the respective clustering algorithms.

Therefore, further research is required to develop accu-

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement 683253/GraphInt) and in part by the Swiss National Science Foundation under grant number CRSII2 147609.

racy metrics that are applicable to overlapping (and multi-resolution) clusterings, that satisfy tight performance constraints and, ideally, that are compatible with the results of standard accuracy metrics used for non-overlapping clustering. Finally, producing a single, easy to interpret value for the final clustering is of importance also, in order to help the user pick the most suitable clustering for a particular use-case and for potentially several accuracy metrics. This issue has been tackled through the formal constraints, introduced for example in [3], [6], and is further discussed in Section III.

To the best of our knowledge, this is the first work discussing all state-of-the-art accuracy metrics applicable to overlapping clustering evaluation on large datasets (i.e., with more than 10^7 elements). Being able to evaluate metrics on large datasets means—in our context—that the evaluation process should be at most: *a) quadratic* in terms of runtime complexity and *b) quasilinear* in terms of memory complexity with the number of elements considered. In addition, we also introduce in this paper a novel indexing technique to reduce the runtime complexity of the Mean F1 score (and a similar metric, NVD [9]) evaluation from $O(N \cdot (|C| + |C'|))$ to $O(N)$, where N is the number of elements in the processed clusters, $|C|$ is the number of resulting clusters and $|C'|$ is the number of ground-truth clusters. Finally, we propose extensions to the state-of-the-art accuracy metrics to satisfy more formal constraints (i.e., to improve their effectiveness) without sacrificing their efficiency. Efficient C++ implementations of *a)* all discussed accuracy metrics and *b)* their improved versions are freely available online as open source utilities as listed in the sections devoted to each metric. All our accuracy metrics are also integrated into an open source benchmarking framework for clustering algorithms evaluation, Clubmark¹ [13], besides being available as dedicated applications.

II. RELATED WORK

Related efforts can be categorized into three main groups, namely: *a)* accuracy metrics for overlapping clustering evaluation (satisfying the complexity constraints mentioned above), *b)* frameworks providing efficient implementations for accuracy evaluation and *c)* formal constraints for accuracy metrics (discussed in Section III).

a) Accuracy Metrics for Overlapping Clusters: The *Omega Index* [4] is the first accuracy metric that was proposed for overlapping clustering evaluation. It belongs to the family of *Pair Counting Based Metrics*. It is a fuzzy version of the Adjusted Rand Index (ARI) [24] and is identical to the Fuzzy Rand Index [25]. We describe the Omega Index in Section IV-A.

Versions of *Normalized Mutual Information (NMI)* suitable for overlapping clustering evaluation were introduced as Overlapping NMI (ONMI)² [19] and Generalized NMI (GNMI) [26] and belong to the family of *Information Theory Based Metrics*. The authors of ONMI suggested to extend Mutual Information

with approximations (introduced in [27]) to find the best matches for each cluster of a pair of overlapping clusterings. This approach allows to compare overlapping clusters, but unlike GNMI we introduce in Section IV-C, it yields values that are incompatible with standard NMI [5] results.

The *Average F1 score* is introduced in [7], [28] and a similar metric, NVD, is introduced in [9]. The Average F1 score belongs to the family of *Cluster Matching Based Metrics* and is described in Section IV-B.

b) Accuracy Measurement Frameworks: A toolkit³ for the parallel measurement of the quality of non-overlapping clusterings on both distributed and shared memory machines is introduced in [9]. This toolkit performs the evaluation of several accuracy metrics (Average F1 score, NMI, ARI and JI) as well as some intrinsic quality metrics, and provides highly optimized parallel implementations of these metrics in C++ leveraging MPI (the Message Passing Interface) and Pthreads (POSIX Threads). Among its accuracy metrics, only Average F1 score is applicable to overlapping clusterings.

WebOCD [10] is an open-source RESTful web framework for the development, evaluation and analysis of *overlapping* community detection (clustering) algorithms. It comprises several baseline algorithms, evaluation metrics and preprocessing utilities. However, since WebOCD (including all its accuracy metrics) is implemented in pure Java as a monolithic framework, many existing implementations of evaluation metrics cannot be easily integrated into WebOCD without either being reimplemented in Java or modifying the framework architecture. A reimplementation of existing metrics is not always possible without a significant performance drop (especially when linking native, high-performance libraries such as Intel TBB, STL or Boost) and time investment.

CoDAR [11] is a framework for community detection algorithm evaluation and recommendation providing user-friendly interfaces and visualizations. Based on this framework, the authors also introduced a study of *non-overlapping* community detection algorithms on *unweighed undirected* networks [14]. Unfortunately, the framework URL provided in the paper refers to a forbidden page, i.e. the implementation is not available to the public anymore.

Circulo [12] is a framework for community detection algorithms evaluation. It executes the algorithms on preliminary uploaded input networks and then evaluates the results using several accuracy metrics and multiple intrinsic metrics.

III. FORMAL CONSTRAINTS ON CLUSTERING EVALUATION METRICS

Four formal constraints for the accuracy metrics were introduced in [3] and shed light on which aspects of the quality of a clustering are captured by different metrics. Two of these constraints (Homogeneity, see Fig. 1 and Completeness, see Fig. 2) were originally proposed in [6] while the other two (Rag Bag, Fig. 3 and Size vs Quantity, Fig. 4) were newly introduced. Besides being intuitive, these four constraints were

¹<https://github.com/eXascaleInfolab/clubmark>

²<https://github.com/eXascaleInfolab/OvpNMI>

³<https://github.com/chenmingming/ParallelComMetric>




Fig. 1. Homogeneity. Clusters should not mix elements belonging to different categories.




Fig. 2. Completeness. Elements belonging to the same category should be clustered together.




Fig. 3. Rag Bag. Elements with low relevance to the categories (e.g., noise) should be preferably assigned to the less homogeneous clusters (macro-scale, low-resolution, coarse-grained or top-level clusters in a hierarchy).




Fig. 4. Size vs Quality (Micro Weighting). A small assignment error in a large cluster is preferable to a large number of similar errors in small clusters.

developed to be formally provable and to clarify the limitations of each metric. We list these constraints in the way they were presented in [3] and later use them to discuss various accuracy metrics. Each constraint is written as an inequality applied to a pair of clusterings, where a quality metric Q (some accuracy metric in our case) from the clustering on the right-hand side is assumed to be better than the one from the left-hand side. Ground-truth clusters are called *categories* for short in the following.

IV. ACCURACY METRICS FOR OVERLAPPING CLUSTERING

Accuracy metrics indicate how much one clustering (i.e., set of clusters) is similar to another (ground-truth) clustering. For each presented metric, we first give its original definition before proposing our extensions and optimizations. Then, we empirically evaluate the aforementioned four formal constraints on samples from [3] and given in Fig. 5-8 denoting the left clustering as *Low* and the right clustering as *High* for each sample. The constraints that are satisfied are marked in italics in the results tables for each discussed metric in the respective subsection. Cluster elements belonging to the same category (ground-truth cluster) in these figures are colored

with the same color and texture, while the formed clusters (results) are shown with oval shapes.




Fig. 5. Homogeneity sample.




Fig. 6. Completeness sample.




Fig. 7. Rag Bag sample.




Fig. 8. Size vs Quality sample.

A. Omega Index

1) *Preliminaries*: The Omega Index [4] is an ARI [24] generalization applicable to overlapping clusters. It is based on counting the number of pairs of elements occurring in exactly the same number of clusters as in the number of categories and adjusted to the expected number of such pairs. Formally, given the ground-truth clustering C' consisting of categories $c'_i \in C'$ and formed clusters $c_i \in C$:

$$\text{Omega}(C', C) = \frac{\text{Obs}(C', C) - \text{Exp}(C', C)}{1 - \text{Exp}(C', C)}. \quad (1)$$

The observed agreement is:

$$\text{Obs}(C', C) = \sum_{j=0}^{\min(J', J)} A_j / P, \quad (2)$$

where J' (J) is the maximal number of categories (clusters) in which a pair of elements occurred, A_j is the number of pairs of elements occurring in exactly j categories and exactly j clusters, and $P = N \cdot (N - 1) / 2$ is the total number of pairs given a total of N elements (nodes of the network being clustered).

The expected agreement is:

$$\text{Exp}(C', C) = \sum_{j=0}^{\min(J', J)} P'_j P_j / P^2, \quad (3)$$

where P'_j (P_j) is the total number of pairs of elements assigned to exactly j categories (clusters).

2) *Proposed Extension (Soft Omega Index)*: The Omega Index evaluates overlapping clusterings by counting the number of pairs of elements present in exactly the same number of clusters as in the number of categories, which does not take into account pairs present in slightly different number of clusters. We propose to fix this issue by normalizing smaller number of occurrences of each pair of elements in all clusters of one clustering by the larger number of occurrences in

another clustering as outlined on line 9 of Algorithm 1. The input data consists of two clusterings (grs, cls), and the $rels$ hashmap relating the clusters to their elements (nodes) for each clustering. The updated computation of the observed agreement of pairs requires also to correct the expected agreement, which is performed on line 19.

Algorithm 1 Soft Omega Index

```

1: procedure OMEGASOFT( $rels, grs, cls$ )
2:    $nobs = 0$  ▷ Observed number
3:    $ngs = \text{vector}(\text{size}(grs))$  ▷ Ranked pairs counts for grs
4:    $ncs = \text{vector}(\text{size}(cls))$  ▷ Ranked pairs counts for cls
5:   for  $ir$  in  $\text{range}(\text{begin}(rels), \text{end}(rels))$  do
6:     for  $jr$  in  $\text{range}(++ir, \text{end}(rels))$  do
7:        $gn = \text{mutualnum}(grs(ir), grs(jr))$ 
8:        $cn = \text{mutualnum}(cls(ir), cls(jr))$ 
9:        $nobs += 1$  if  $gn == cn$  else  $\min(gn, cn) /$ 
max( $gn, cn$ )
10:       $++ngs[gn]; ++ncs[cn]$ 
11:     end for
12:   end for
13:    $nexp = 0; szmin = \min(\text{size}(ngs), \text{size}(ncs))$ 
14:   for  $i$  in  $\text{range}(szmin)$  do
15:      $nexp += ngs[i] \cdot ncs[i]$ 
16:   end for
17:    $rns = ngs$  if  $\text{size}(ngs) > szmin$  else  $ncs$ 
18:   for  $i$  in  $\text{range}(szmin, \text{size}(rns))$  do
19:      $nexp += rns[i]$ 
20:   end for
21:    $nps = \text{size}(\text{nodes}(rels)); nps *= (nps - 1) / 2$ 
22:    $nexp /= nps$ 
23:   return  $(nobs - nexp) / (nps - nexp)$ 
24: end procedure

```

Thus, *OmegaSoft* has the same definition as Eq. 1, except the observed agreement number is evaluated as:

$$Obs_{soft}(C', C) = \sum_{j=0}^{\max(J', J)} Anorm_j / P, \quad (4)$$

where $Anorm_j$ is the number of pairs of elements occurring in exactly j' and j clusters of the clusterings and being weighted by $\min(j', j) / \max(j', j)$.

The expected agreement is:

$$Exp_{soft}(C', C) = \left(\sum_{j=0}^{Jmin} P_j' P_j + \sum_{j=Jmin+1}^{\max(J', J)} Prem_j \right) / P^2, \quad (5)$$

where $Jmin = \min(J', J)$, $Prem_j = P_j$ if $\min(J', J) = J'$ and $Prem_j = P_j'$ otherwise.

Note that for non-overlapping clusterings (i.e., when the membership in clusters equals to 1 $\implies J' = J = 1$), the Soft Omega Index is equivalent to the original Omega Index, which is equivalent to ARI.

3) *Evaluation and Constraints Matching*: A counterexample outlining the issue of the Omega Index when discarding partially matching pairs of elements is shown in Table I, where each of the four categories (C1'-C4') consists of 3 elements and the total number of elements is 4 (#1-#4). The first clustering algorithm has a *Low* accuracy and discovers only two clusters as shown in Table I. The second clustering algorithm (*High*) performs much better discovering all four clusters but not all elements of the respective categories as shown in Table I. The original Omega Index fails to discriminate these cases yielding 0 for both cases, whereas the Soft Omega Index clearly differentiates the more accurate solution.

TABLE I
ACCURACY EVALUATION OF *Low* AND *High* VS *Ground-truth* CLUSTERINGS BY OMEGA INDEX AND SOFT OMEGA INDEX.

Metrics \ Clusterings	Ground-truth	
	Low	High
	C1': 1 2 3	
	C2': 2 3 4	
	C3': 3 4 1	
	C4': 4 1 2	
Omega Index	0	0
Soft Omega Index	0	0.33

The empirical satisfaction of the formal constraints for both versions of Omega Index is given in Table II and discussed below in Section IV-D. The computational complexity $O(N^2)$ and the memory complexity is $O(|C| + |C'|) \approx^4 O(\sqrt{N})$ for both implementations, where N is the number of elements in the clustering. Implementations of both the original and Soft Omega Index are provided in the open source *xmeasures*⁵ utility and are available for free.

TABLE II
FORMAL CONSTRAINTS FOR SOFT AND ORIGINAL OMEGA INDEX.

Clusterings Metrics	Homogen.		Comple.		RagBag		SzQual.	
	low	high	low	high	low	high	low	high
[Soft] Omega Index	0.247	0.282	0.244	0.311	0.4	0.4	0.804	0.804

B. Mean F1 Score

1) *Preliminaries*: The *Average F1 score (F1a)* is a commonly used metric to measure the accuracy of clustering algorithms [7], [28], [29]. F1a is defined as the average of the weighted F1 scores [30] of a) the *best matching* ground-truth clusters to the formed clusters and b) the *best matching* formed clusters to the ground-truth clusters. Formally, given the ground-truth clustering C' consisting of clusters $c'_i \in C'$

⁴On average, the number of clusters in most real-world networks consisting of N nodes is \sqrt{N}

⁵<https://github.com/eXascaleInfolab/xmeasures>

(called categories) and clusters $c_i \in C$ formed by the evaluating clustering algorithm:

$$F1a(C', C) = \frac{1}{2}(F_{C',C} + F_{C,C'}), \quad (6)$$

where

$$F_{X,Y} = \frac{1}{|X|} \sum_{x_i \in X} F1(x_i, g(x_i, Y)),$$

$$g(x, Y) = \{\text{argmax}_y F1(x, y) \mid y \in Y\}, \quad (7)$$

where $F1(x, y)$ is the F1 score of the respective clusters.

2) *Proposed Extensions (F1h and F1p)*: The F1a definition yields *non-indicative* values of $F1a \in [0, 0.5]$ when evaluating a large number of clusters. In particular, for clusters formed by taking all possible combinations of the nodes, $F1a > 0.5$ ($F1_{C',C} = 1$ since for each category there exists the exactly matching cluster, $F1_{C,C'} \rightarrow 0$ since majority of clusters have low similarity to the categories). To address this issue, we suggest to use the harmonic mean instead of the arithmetic mean (average). We introduce the *harmonic F1 score (F1h)* as:

$$F1h(C', C) = \frac{2F_{C',C}F_{C,C'}}{F_{C',C} + F_{C,C'}}. \quad (8)$$

$F1h \leq F1a$ since the harmonic mean cannot be larger than the arithmetic mean. In our case, $F_{C',C}$ can be interpreted as a recall and $F_{C,C'}$ as a precision of the evaluating clustering C and the ground-truth clustering C' .

$F1h$ is more indicative than $F1a$ but both measures do not satisfy the Homogeneity constraint, as they penalize local best matches too severely. We propose to evaluate the probability of the local best matches rather than the F1 score to address the outlined issue. Our new metric $F1p$ is the harmonic mean (i.e. F1 measure) of the average over each clustering of the best local probabilities for each cluster. $F1p$ corresponds to the expected probability of the best match of the clusterings unlike $F1h$, which corresponds to the expected worst-case of the best match in the clusterings. Formally, $F1p$ is evaluated similarly to $F1h$, except that the local matching function $pprob$ given in Eq. 9 replaces $f1$ given in Eq. 10.

$$pprob(m, c', c) = m/|c'| * m/|c| = \frac{m^2}{|c'| * |c|}, \quad (9)$$

$$f1(m, c', c) = 2 \frac{m/|c'| * m/|c|}{m/|c'| + m/|c|} = 2 \frac{m}{|c'| + |c|}, \quad (10)$$

where m is the *contribution* of matched elements between the cluster c and category c' . The notations of contribution and $|x|$ (size of the cluster x) vary for the overlapping and multi-resolution clusterings and are discussed in Section IV-D. For multi-resolution and non-overlapping clusterings, $|x|$ is simply equal to the number of elements in the cluster x , and the contribution of each element is equal to 1. For overlapping clusterings, $|x|$ is equal to the total contribution of elements in cluster x , where each element x_i contributes the value $1/\text{shares}(x_i)$ given that x_i is a member of the number $\text{shares}(x_i)$ of (overlapping) clusters.

3) *Optimizations (Efficient Indexing) for the F1 Metric Family*: We propose an efficient indexing technique to reduce the computational complexity of computing the *Mean F1 score* metric family ($F1a, F1h$ and $F1p$). Our technique is based on dedicated data structures described below. When loading the clusterings, we create a *rels* hashmap relating the clusters to their elements (nodes) for each clustering. Besides the member nodes, our cluster data structure holds also: a) an accumulator *ctr* for the matching contributions of the member nodes together with the pointer to the originating cluster from which these contributions are formed, and b) the local contributions *cont* for all members nodes of the cluster. This data structure allows to evaluate the metrics using a single pass over all members of all clusters. The content of the *ctr* attribute is reset on line 9 in Algorithm 2 when adding a value of the contribution together with a distinct cluster pointer from the already stored one. The main procedure to evaluate the aforementioned best matches for the clustering *cls* is listed in Algorithm 2 considering a) the `fmatch` matching function (`f1` or `pprob` given in Eq. 9-10) parameterized as *prob* argument, and b) the overlapping or multi-scale clusters evaluation semantics parameterized as *ovp*.

Algorithm 2 Best matches evaluation using efficient indexing technique

```

1: procedure BMATCHES(grels, crels, cls, prob, ovp)
2:   bms = vector()
3:   fmatch = pprob if prob else f1  ▷ Matching function
4:   for c in cls do  ▷ Traverse each cluster
5:     bmt = 0  ▷ Value of the best match
6:     for nd in members(c) do
7:       ndcls = grels[nd]  ▷ Node clusters
8:       for cn in ndcls do  ▷ Evaluate node clusters
9:         cn.ctr.add(c, 1 / max(size(grels[nd]),
size(crels[nd])) if ovp else 1)
10:        mt = fmatch(cn.ctr, cn.cont, c.cont)
11:        if bmt < mt then
12:          bmt = mt
13:        end if
14:      end for
15:    end for
16:    if prob then
17:      bmt =  $\sqrt{\text{bmt}}$ 
18:    end if
19:    bms.append(bmt)
20:  end for
21:  return bms
22: end procedure

```

4) *Evaluation and Constraints Matching*: Our new indexing technique reduces the computational complexity of Mean F1 Score evaluation from $O(N \cdot (|C| + |C'|))$ [9] $\simeq^4 O(N\sqrt{N})$ to $O(N \cdot s) \simeq O(N)$, where N is the number of elements in the processing clusters and the constant s is the average membership of the elements, which typically is $\in [1, 2]$ for overlapping clusterings. Implementations of all *MF1* metrics

($F1a$, $F1h$, $F1p$) are provided in the open source *xmeasures*⁵ utility and are available for free. The empirical satisfaction of the formal constraints for all $MF1$ metrics is given in Table III and discussed in Section IV-D.

TABLE III
FORMAL CONSTRAINTS FOR MF1 METRIC FAMILY.

Clusterings Metrics	Homogen.		Comple.		RagBag		SzQual.	
	low	high	low	high	low	high	low	high
F1a	0.646	0.646	0.639	0.663	0.641	0.630	0.795	0.936
F1h	0.646	0.646	0.639	0.660	0.639	0.630	0.795	0.935
F1p	0.665	0.672	0.686	0.703	0.693	0.693	0.819	0.942

C. Generalized NMI

1) *Preliminaries*: Generalized NMI (GNMI)⁶ [26] uses a stochastic process to compare overlapping clusterings and feeds the random variables of the process into the standard definition of mutual information (MI) [31]. MI is evaluated by taking all pairs of clusters from the formed and ground-truth clusterings respectively and counts the number of common elements in each pair. Formally, given the ground-truth clustering C' consisting of clusters $c' \in C'$ and the formed clusters $c \in C$, mutual information is defined as:

$$I(C' : C) = \sum_{c' \in C'} \sum_{c \in C} p(c', c) \log_2 \frac{p(c', c)}{p(c')p(c)}, \quad (11)$$

where $p(c', c)$ is the normalized number of common elements in the pair of (*category*, *cluster*), $p(c')$ and $p(c)$ is the normalized number of elements in the categories and formed clusters respectively. The normalization is performed using the total number of elements in the clustering, i.e. the number of nodes in the input network.

Normalized Mutual information (NMI) [5] performs normalization of MI by *maximum value*, *arithmetic or geometric mean* of the unconditional entropies $H(C')$ and $H(C)$ of the clusterings. Normalization by the maximum value of the entropies is the standard approach, which is also considered as the most discriminative one [19], [26]:

$$NMI(C', C) = \frac{I(C' : C)}{\max(H(C'), H(C))}. \quad (12)$$

The unconditional entropy $H(X)$ [32] of clusters $x \in X$ is:

$$H(X) = - \sum_{x \in X} p(x) \log_2 p(x). \quad (13)$$

2) *Proposed Extension and Optimizations*: The GNMI approach of using a stochastic process provides the only known way to evaluate *standard NMI* for overlapping clusterings. The original GNMI implementation⁶ uses Intel's TBB library (lightweight threads) to execute the stochastic process on all available CPUs efficiently. However, the original implementation has several shortcomings, which makes it inapplicable to large datasets:

- its hard-coded maximal number of stochastic events (successful samples) `EVCOUNT_THRESHOLD` is not adequate for handling both small and large datasets. Moreover, the original value is too small for large datasets;
- its fully random sampling of the cluster elements has a too high computational complexity on large datasets to produce results considering a reasonably small evaluation error (default value is 0.01);
- the two aforementioned points cause significant errors on small datasets and loss of convergence on large datasets while consuming significant computational resources.

We optimized and extended the original version addressing these issues in the following ways. First, instead of the hard-coded `EVCOUNT_THRESHOLD`, we dynamically evaluate the maximal number of stochastic events as:

$$evsmax = \max(\min(mbs', mbs), \frac{1}{rerr\sqrt{rrisk}}), \quad (14)$$

where mbs' and mbs are the total membership of the elements in the categories and clusters respectively, $rerr$ is the admissible error and $rrisk$ is the complement of the resulting confidence; $rerr$ and $rrisk$ are specified as input arguments with a default value equal to 0.01.

Second, we extended the original `try_get_sample` procedure with the weighed adaptive sampling given in Algorithm 3. The original version randomly takes the first node (cluster element) among all nodes in the clusterings as shown on line 2 and then applies *mixer* until it returns `false` or all nodes are traversed in a randomized order. We traverse the nodes located in the same cluster of the formed (c) or ground-truth (g) clusterings on line 7 and weight the shared nodes inversely to their membership. The weighting is performed to discount the contribution (importance) of frequent nodes compared to rare nodes since we traverse only a fraction of all nodes and they may have varying membership ≥ 1 . Indexes on the matched ground-truth and formed clusters are stored in the *mixer* and their matching probability is returned explicitly as *importance*.

Algorithm 3 Adaptive sampling of elements in GNMI

```

1: procedure TRY_GET_SAMPLE(nodes, rrisk, mixer)
2:   node = nodes[rand(size(nodes))]
3:   g, c = clsPair(node)
4:   attempts = (size(g) + size(c)) / (2 · rrisk)
5:   importance = 1 / max(√size(g) · size(c), 1)
6:   adone = 1
7:   while mixer.apply(g, c) and ++adone ≤ attempts do
8:     cm = rand(g, c)
9:     ndm = cm[rand(size(cm))]
10:    g, c = clsPair(ndm)
11:    importance += 1 / max(√size(g) · size(c), 1)
12:   end while
13:   return importance / adone
14: end procedure

```

In addition, we performed some technical optimizations to

⁶<https://github.com/eXascaleInfolab/GenConvMI>

reduce the number of memory allocations and copies, calls to external functions, etc. Our *extended implementation of GNMI*⁷ is open source and available for free.

3) *Evaluation and Constraints Matching*: The empirical satisfaction of the formal constraints for *NMI*, the original *GNMI_{orig}* and our *GNMI* implementations is given in Table IV and discussed below in Section IV-D. Since GNMI-

TABLE IV
FORMAL CONSTRAINTS FOR NMI, ORIGINAL GNMI AND OUR GNMI.

Clusterings Metrics	Homogen.		Comple.		RagBag		SzQual.	
	low	high	low	high	low	high	low	high
NMI	0.450	0.555	0.546	0.546	0.434	0.434	0.781	0.888
GNMI _{orig}	0.512	0.598	0.572	0.632	0.417	0.397	0.808	0.877
GNMI	0.448	0.557	0.546	0.547	0.434	0.436	0.781	0.888

s yields stochastic results, we report the median value over 5 runs with the same default values of the error and risk arguments. As the table clearly shows, our GNMI implementation is much more accurate than the original one and yields values equal to the original NMI within the specified admissible error (0.01) on non-overlapping clusterings. The empirical evaluation on *a*) the synthetic datasets formed using the LFR⁸ [33] framework and *b*) the large real-world networks with ground-truth⁹ introduced in [29] show that our implementation is one order of magnitude faster on datasets with 10^4 nodes and two orders of magnitude faster on datasets with 10^6 nodes than the original GNMI implementation. The actual computational complexity of both GNMI implementations depends on the structure of the clusters and on the number of overlaps in the clusterings. Moreover, for very dissimilar clusterings the evaluation might not converge at all, which is a disadvantage of the stochastic evaluation. The worst case computational complexity for our GNMI implementation is $O(N \cdot s \cdot |\bar{c}| \cdot |C|) \approx O(N \cdot |\bar{c}| \cdot |C|)$ while it is $O(\text{EVCOUNT_THRESHOLD} \cdot N \cdot |C|)$ for the original GNMI, where N is the number of elements in the evaluating clusterings, $|C|$ is the number of clusters and $|\bar{c}|$ is the average size of the clusters.

D. Discussion

When evaluating clusterings, it is important to *a*) distinguish overlapping from multi-resolution (including hierarchical) clusterings and *b*) consider the limitations of the applied accuracy metric in order to produce meaningful results. First, we discuss the differences when handling overlapping clusterings versus multi-resolution clusterings having non-overlapping clusters on each resolution and how they affect accuracy evaluations.

In the case of overlapping clusterings, a node x_i can be shared between s clusters and has equal membership in each of them (e.g., a person may spend time for several distinct hobbies but the more hobbies are involved the less time the person devotes to each one). Thus, the membership

contribution x_i to each of the clusters is $1/s$. In case of non-overlapping, multi-resolution clusterings, a node x_i may be a member of s nested clusters taken at different resolutions (granularities), but here x_i fully belongs to each of these clusters having a membership contribution equal to 1 (e.g., a student tacking a course can be a full member of the course, as well as of the department offering the course and of the whole university). These distinct semantics for the elements shared in a clustering are represented by the *ovp* argument in our *xmeasures*⁵ utility for all MF1 metrics.

We summarize the advantages and limitations of the various metrics discussed in this paper as follows:

Omega Index pros: its values are not affected by the number of clusters (unlike NMI) and have an intuitive interpretation (0 means equal quality to a random clustering).

Omega Index cons: it performs purely when applied to multi-resolution clusterings and has the highest computational complexity among all considered metrics ($O(N^2)$).

MF1 pros: it has the lowest computational complexity (linear when using our indexing technique).

MF1 cons: it evaluates the best-matching clusters only, which gives an unfair advantage to the clusterings with larger numbers of clusters (which is partially addressed by the application of the harmonic mean in *F1h* and *F1p* instead of the average).

GNMI pros: it parallelizes very well and inherits NMI’s pros, namely it evaluates the full matching of all clusters (unlike MF1). Also, it is well-grounded formally in Information Theory.

GNMI cons: the convergence in the stochastic implementation is not guaranteed (though loss of convergence typically indicates a relevance close to zero); the stochastic process yields non-deterministic results and is computationally heavy; its execution time is hard to estimate. In addition, GNMI inherits the NMI’s cons, namely the results depends on the number of clusters being evaluated and increase up to ≈ 0.3 for large numbers of clusters.

TABLE V
FORMAL CONSTRAINTS FOR OMEGA INDEX, MF1 AND GNMI.

Metrics\Clusterings	Homogen.	Comple.	RagBag	SzQual.
[Soft] Omega Index	+	+		
F1h		+		+
F1p	+	+		+
GNMI	+			+

Empirical evaluation of the formal constraints satisfaction for all discussed metrics on the original intuitive samples from [3] is given in Table V. As it shown in the table, none of the metrics satisfies the RagBag constraint, which is essential in practice when evaluating the multi-resolution or hierarchical clusterings since the more fine-grained (higher resolutions) are expected to have less noisy structure. Otherwise, the proposed *F1p* metric performs the best according to the empirical satisfaction of the formal constraints having the lowest computational complexity (linear on the number of

⁷<https://github.com/eXascaleInfolab/GenConvNMI>

⁸https://github.com/eXascaleInfolab/LFR-Benchmark_UndirWeightOvp

⁹<https://snap.stanford.edu/data/#communities>

elements being clustered) compared to GNMI and Omega Index.

V. CONCLUSIONS

In this paper, we discussed the state-of-the-art accuracy metrics applicable to overlapping clustering evaluations on large datasets and introduced several optimizations and extensions of the discussed metrics. In particular, we introduced an efficient indexing technique to speedup the evaluation of Mean F1 score from $O(N\sqrt{N})$ to $O(N)$, where N is the number of elements in the clustering. We proposed an adaptive sampling strategy for *GNMI*, which not only speeds up the evaluation by orders of magnitude but also improves the precision of the metric. In addition, we proposed two extensions of the Average F1 score ($F1a$), namely $F1h$ and $F1p$. $F1h$ addresses the issues of the loss of indicativity of $F1a$ in the range $[0, 0.5]$ while $F1p$ empirically improves the satisfaction of the formal constraints we consider without sacrificing efficiency. We also proposed an extension of the Omega Index called *Soft Omega Index*, which is equivalent to the original Omega Index evaluating non-overlapping clusterings and yields more discriminative results for overlapping clusterings due to the fact that it considers partial matches of pairs of elements.

Besides the proposed optimizations and extensions of the accuracy metrics, we discussed formal constraints for each metric, as well as their applicability and limitations for specific cases in overlapping and multi-resolutions clusterings. Our analysis of the metrics should provide insight for their future usage and should hopefully help identify the best performing clustering algorithm for particular user's needs or tasks.

We freely provide implementations of all discussed metrics, and are also integrating them into an open source benchmarking framework for clustering algorithms evaluation, Clubmark¹.

REFERENCES

- [1] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E*, vol. 69, no. 2, p. 026113, 2004.
- [2] R. Kannan, S. Vempala, and A. Vetta, "On clusterings: Good, bad and spectral," *J. ACM*, vol. 51, no. 3, pp. 497–515, May 2004.
- [3] E. Amigó, J. Gonzalo, J. Artilles, and F. Verdejo, "A comparison of extrinsic clustering evaluation metrics based on formal constraints," *Inf. Retr.*, vol. 12, no. 4, Aug. 2009.
- [4] L. M. Collins and C. W. Dent, "Omega: A general formulation of the rand index of cluster recovery suitable for non-disjoint solutions," *Multivariate Behavioral Research*, vol. 23, no. 2, pp. 231–242, 1988.
- [5] L. Danon, A. Díaz-Guilera, J. Duch, and A. Arenas, "Comparing community structure identification," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 9, p. 8, Sep. 2005.
- [6] A. Rosenberg and J. Hirschberg, "V-measure: A conditional entropy-based external cluster evaluation measure," ser. EMNLP-CoNLL'07, pp. 410–420.
- [7] J. Yang and J. Leskovec, "Overlapping community detection at scale: A nonnegative matrix factorization approach," ser. WSDM '13. ACM, pp. 587–596.
- [8] H. Rosales-Méndez and Y. Ramírez-Cruz, "Cice-bcubed: A new evaluation measure for overlapping clustering algorithms," ser. CIARP'13.
- [9] M. Chen, S. Liu, and B. K. Szymanski, "Parallel toolkit for measuring the quality of network community structure," in *ENIC*, 2014, pp. 22–29.
- [10] M. Shahriari, S. Krott, and R. Klamma, "Webcod: A restful web-based overlapping community detection framework," ser. i-KNOW '15. ACM.
- [11] X. Ying, C. Wang, M. Wang, J. X. Yu, and J. Zhang, "Codar: Revealing the generalized procedure & recommending algorithms of community detection," ser. SIGMOD '16. ACM, 2016, pp. 2181–2184.
- [12] P. Mazzuca and Y. T., "Circulo: A Community Detection Evaluation Framework," Feb. 2015. [Online]. Available: <https://www.lab41.org/circulo-a-community-detection-evaluation-framework/>
- [13] A. Lutov, M. Khayati, and P. Cudré-Mauroux, "Clubmark: a parallel isolation framework for benchmarking and profiling clustering algorithms on numa architectures," in *ICDMW*, 2018, pp. 1481–1486.
- [14] M. Wang, C. Wang, J. X. Yu, and J. Zhang, "Community detection in social networks: An in-depth benchmarking study with a procedure-oriented framework," *Proc. VLDB Endow.*, vol. 8, pp. 998–1009, Jun. 2015.
- [15] S. Harenberg, G. Bello, L. Gjøltema, S. Ranshous, J. Harlalka, R. Seay, K. Padmanabhan, and N. Samatova, "Community detection in large-scale networks: A survey and empirical evaluation," *WIREs Comput. Stat.*, vol. 6, pp. 426–439, Nov. 2014.
- [16] M. Hascoët and G. Artignan, "Evaluation of Clustering Algorithms: a methodology and a case study," Tech. Rep. RR-14008, Sep. 2014.
- [17] W. Saelens, R. Cannoodt, and Y. Saeyns, "A comprehensive evaluation of module detection methods for gene expression data," *Nat. Commun.*, vol. 9, p. 12, 2018.
- [18] V. Melnykov and R. Maitra, "Carp: Software for fishing out good clustering algorithms," *J. Mach. Learn. Res.*, vol. 12, pp. 69–73.
- [19] A. F. McDaid, D. Greene, and N. J. Hurley, "Normalized mutual information to evaluate overlapping community finding algorithms," *CoRR*, vol. abs/1110.2515, 2011.
- [20] S. Zhang, R.-S. Wang, and X.-S. Zhang, "Identification of overlapping community structure in complex networks using fuzzy c-means clustering," *Physica A: Statistical Mechanics and its Applications*, vol. 374, no. 1, pp. 483–490, 2007.
- [21] V. Nicosia, G. Mangioni, V. Carchiolo, and M. Malgeri, "Extending the definition of modularity to directed graphs with overlapping communities," *J. Stat. Mech.*, vol. 3, p. 24, Mar. 2009.
- [22] A. Lázár, D. Abel, and T. Vicsek, "Modularity measure of networks with overlapping communities," *EPL (Europhysics Letters)*, vol. 90, no. 1, p. 18001, 2010.
- [23] M. Chen and B. K. Szymanski, "Fuzzy overlapping community quality metrics," *Social Network Analysis and Mining*, vol. 5, no. 1, p. 40, Jul 2015.
- [24] L. Hubert and P. Arabie, "Comparing partitions," *Journal of Classification*, vol. 2, no. 1, pp. 193–218, Dec 1985.
- [25] E. Hullermeier and M. Rifqi, "A fuzzy variant of the rand index for comparing clustering structures," ser. IFSA-EUSFLAT 2009.
- [26] A. V. Esquivel and M. Rosvall, "Comparing network covers using mutual information," *CoRR*, vol. abs/1202.0425, 2012.
- [27] A. Lancichinetti, S. Fortunato, and J. Kertész, "Detecting the overlapping and hierarchical community structure in complex networks," *New J. Phys.*, vol. 11, no. 3, p. 033015, 2009.
- [28] A. Prat-Pérez, D. Dominguez-Sal, and J.-L. Larriba-Pey, "High quality, scalable and parallel community detection for large real graphs," ser. WWW'14. ACM, pp. 225–236.
- [29] J. Yang and J. Leskovec, "Defining and evaluating network communities based on ground-truth," *Knowl. Inf. Syst.*, vol. 42, pp. 181–213, Jan. 2015.
- [30] C. J. V. Rijsbergen, *Information Retrieval*, 2nd ed. Newton, MA, USA: Butterworth-Heinemann, 1979.
- [31] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 1951.
- [32] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, July 1948.
- [33] A. Lancichinetti and S. Fortunato, "Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities," *Phys. Rev. E*, vol. 80, p. 016118, Jul 2009.