

A Pragmatic Method for Comparing Clusterings with Overlaps and Outliers

Ryan DeWolfe*, Paweł Prałat†, François Thériberge‡

Abstract

Clustering algorithms are an essential part of the unsupervised data science ecosystem, and extrinsic evaluation of clustering algorithms requires a method for comparing the detected clustering to a ground truth clustering. In a general setting, the detected and ground truth clusterings may have outliers (objects belonging to no cluster), overlapping clusters (objects may belong to more than one cluster), or both, but methods for comparing these clusterings are currently undeveloped. In this note, we define a pragmatic similarity measure for comparing clusterings with overlaps and outliers, show that it has several desirable properties, and experimentally confirm that it is not subject to several common biases afflicting other clustering comparison measures.

1 Introduction

Clustering [13] (or community detection [8]) is one of the fundamental problems in unsupervised machine learning, and typically involves finding groups of “similar” data within a dataset. Many clustering algorithms often only consider partitions of the data (a clustering where each object belongs to exactly one cluster), but several recent algorithms allow for outliers [3, 23], overlapping clusters [28], or both [6, 20]. One method to evaluate the performance of a clustering algorithm is extrinsic clustering validity: the algorithm is judged by its ability to find a known ground truth clustering in the data. Thus, extrinsic validity benchmarks require a method to compare clusterings [11].

In this note, we use set-matching [25] and the F^* score [12] to define a pragmatic method for comparing clusterings that could have overlaps or outliers. The method is simple, computationally efficient, and exhibits several intuitive properties one would expect in a comparison measure.

*Department of Mathematics, Toronto Metropolitan University, Toronto, Canada; ryan.dewolfe@torontomu.ca

†Department of Mathematics, Toronto Metropolitan University, Toronto, Canada; pralat@torontomu.ca

‡Tutte Institute for Mathematics and Computing, Ottawa, Canada; theberge@ieee.org

2 Related Work

Comparing partitions is a well-studied problem, and there are a wide variety of measures that are used [11], which can be broadly classified as either pair-counting, information theoretic, or set-matching [25]. In this section, we review some extensions to overlapping clusterings, although to the best of our knowledge, none have explicitly considered outliers.

The *Omega* index [5] extends the pair-counting Adjusted Rand Index (ARI) [14] by extending the contingency table from agree versus disagree to counting the number of clusters in common. This extension is reasonable and is equivalent to the ARI for partitions, but is quite harsh since it does not detect any similarity if a pair of objects does not share the exact same number of clusters in both clusterings. Furthermore, the non-adjusted version of the Omega index suffers from the same issues as the non-adjusted Rand index: values tend to be large, even for random clusterings. The adjustment proposed implicitly assumes a null-model for generating random clusterings, the choice of which has been shown to have an effect for comparing partitions [9], and generally requires careful consideration of the assumptions when extending beyond partitions [7].

Lancichinetti et al. [18] proposed the overlapping Normalized Mutual Information (oNMI), an extension of the information theoretic Normalized Mutual Information [27] using a set matching framework. The oNMI (and NMI) requires an adjustment term to normalize values across datasets, and McDaid et al. [21] note that the original adjustment leads to some unintuitive behaviour, and suggests an alternative. In fact, there are 5 possible adjustment terms and, similar to the previously mentioned null-models, any choice must be carefully considered. In our experiments, we test both suggested adjustments, referring to that of Lancichinetti et al. [18] as oNMI_{LFK} and that of McDaid et al. [21] as oNMI_{MGH} .

Finally, by shifting the focus from clusters to objects, the Element-Centric Similarity (ECS) [10] is a new class of method. The method constructs a *cluster-induced element graph* by placing a directed and non-uniformly weighted clique over each cluster (with weights chosen so that the degree of each object is split evenly across its clusters), and comparing the personalized pagerank vectors of each object. This method is intuitive and allows for interesting object-level agreement analysis, but has several drawbacks for comparing clusterings. Critically, the ECS may assign maximum similarity to non-identical clusterings; for example, a single cluster that contains all objects compared to any uniform Balanced Incomplete Block Design (BIBD) over the objects will achieve maximum similarity. A BIBD is a clustering where each pair of objects appears in the same number of clusters and each cluster has the same size. These symmetries create a cluster induced element graph that is a clique with equal weights, which is identical (up to a constant scaling, which is ignored by the personalized pagerank computation) to the cluster-induced element graph of a single cluster containing all objects. Furthermore, the ECS is computationally difficult for large datasets with overlapping clusters, and the natural extension would treat outliers the same as clusters of size one.

Finally, we briefly mention the *problem of matching* [22] that affects set-matching measures [25] like the oNMI and our proposed measure. The problem of matching refers to the indifference of a measure to the part of each cluster that is unmatched. In particular, if the

unmatched portion of a cluster is concentrated in one other cluster, then it is somehow less disrupted than if the unmatched portion was spread across many clusters. However, while it may be less disrupted in an edit distance sense (and only if moving multiple objects between clusters costs less than moving each object individually), we disagree that intuition dictates that these disruptions should not be equally dissimilar to the original. In fact, by measuring the ability of each cluster to find a match in the other clustering, a set-matching measure provides interesting cluster level information that could be useful for detecting where the clusterings disagree, possibly complementing the object level information available from the ECS [10].

3 The Method

For a set of objects $X = \{x_1, x_2, \dots, x_n\}$, we define a *clustering* as set of non-empty subsets of X . We call each element of a clustering a *cluster*. Note that with this definition, each cluster in a clustering is distinct, and the empty set cannot be a cluster. However, we allow the clustering itself to be the empty set.

To define methods for comparing clusterings $\mathcal{C}^{(1)}$ and $\mathcal{C}^{(2)}$, we first define some methods for comparing two individual clusters. A cluster similarity function is a map between two subsets of X and a value in $[0, 1]$ that measures some notion of similarity (with 1 denoting maximum similarity). The standard precision (p), recall (r), and F_1 similarity functions between two clusters $C_i, C_j \subseteq X$ are defined as:

$$p(C_i, C_j) = \frac{|C_i \cap C_j|}{|C_i|},$$

$$r(C_i, C_j) = \frac{|C_i \cap C_j|}{|C_j|},$$

$$F_1(C_i, C_j) = \frac{2}{1/p(C_i, C_j) + 1/r(C_i, C_j)}.$$

Precision and recall both have nice interpretations as the probability of being in a cluster conditional on being in other. The F_1 score is very popular to combine the precision and recall into a single score, but has gathered criticism for being pragmatic but ultimately unjustified [4]. In [12], the authors argue that another combination of precision and recall, which they call the F^* score, is more interpretable. Furthermore, they show that F^* is a monotonic transformation of F_1 , which means any ranking of similarities using F_1 will be the same when using F^* . The F^* score is equivalent to the well-known Jaccard index for set similarity, and is defined as

$$F^*(C_i, C_j) = \frac{F_1(C_i, C_j)}{2 - F_1(C_i, C_j)} = \frac{|C_i \cap C_j|}{|C_i \cup C_j|}.$$

For the rest of the paper, we will use F^* . It could be replaced by any of the above (or other) set similarity measures, but it is easy to work with and will serve as a general comparison measure.

Next, we define methods for measuring how well each cluster in $\mathcal{C}^{(1)}$ can find a matching cluster in $\mathcal{C}^{(2)}$. For a cluster $C_i \in \mathcal{C}^{(1)}$, we find a matching cluster by taking the maximum similarity when comparing to each cluster in $\mathcal{C}^{(2)}$:

$$F^*(C_i, \mathcal{C}^{(2)}) = \max_{C_j \in \mathcal{C}^{(2)}} \{F^*(C_i, C_j)\}.$$

We define the edge case $F^*(C_i, \emptyset) = 0$ since C_i cannot find a cluster with any similarity.

Then, to summarize the match quality from each cluster, we compute the weighted average similarity.

$$\hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = \sum_{C_i \in \mathcal{C}^{(1)}} \frac{|C_i|}{\sum_{C_k \in \mathcal{C}^{(1)}} |C_k|} F^*(C_i, \mathcal{C}^{(2)}).$$

This function is a measure of how well each cluster in $\mathcal{C}^{(1)}$ can find a match in $\mathcal{C}^{(2)}$, but not necessarily the reverse. This asymmetric similarity measure could be useful in domains where the performance of a clustering algorithm should be measured by its ability to find clusters that actually exist, but where we are not concerned with complete coverage; for example, in some instances of exploratory data analysis.

Of course, we can also find the best match for each cluster from $\mathcal{C}^{(2)}$ in $\mathcal{C}^{(1)}$ with $\hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)})$. Finally, we use a linear combination with equal weights to combine both directions into a single symmetric similarity score;

$$F_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = 0.5\hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) + 0.5\hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}).$$

The similarity measure F_w^* serves as a simple score to evaluate how similar the clusters of two clusterings are, but crucially does not consider objects that are not part of any cluster. We call an object $x_i \in X$ an *outlier* with respect to a clustering \mathcal{C} if $x \notin \bigcup_{C \in \mathcal{C}} C$. For cases where most of the objects are not outliers, F_w^* might be acceptable, but we can add a term for matching the outliers between clusterings. Let $O^{(1)}$ be the set of outliers with respect to $\mathcal{C}^{(1)}$, and likewise for $O^{(2)}$ with $\mathcal{C}^{(2)}$. The outlier aware versions of \hat{F}_w^* and F_w^* are defined as

$$\hat{F}_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = \frac{|O^{(1)}|}{|X|} F^*(O^{(1)}, O^{(2)}) + \frac{|X \setminus O^{(1)}|}{|X|} \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)})$$

and

$$F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = 0.5\hat{F}_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) + 0.5\hat{F}_{wo}^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}).$$

4 Properties

In this section, we analyze some basic properties of the proposed F_{wo}^* score, and establish that it is reasonable in matching the intuition and general use-case of comparing clustering algorithms. First, we state a few properties that are absolutely required for an intuitive similarity measure, the proofs of which are elementary and are omitted for brevity.

Proposition 1. *The proposed measure F_{wo}^* has the following properties.*

1. F_{wo}^* is label invariant.
2. $F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) \in [0, 1]$, and $F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = 1$ if and only if $\mathcal{C}^{(1)} = \mathcal{C}^{(2)}$.
3. F_{wo}^* is symmetric.
4. Single object clusters are not the same as outliers.

The first property is the essential difference between clustering and multi-class classification tasks. Implementations generally assign a label to each cluster (often using the natural numbers 1 to $|\mathcal{C}|$), but such a labelling is arbitrary and should not affect the similarity to another clustering. The second property is called *normalization* [27], and it is defined as a bounded range for the measure independent of the particular clusterings under consideration. A normalized measure means that the output is on the same scale regardless of the clusterings, and scores are comparable even if the clusterings have a different number of clusters or objects. The final property is the key differentiation of our method with those presented in Section 2, in the sense that the obvious extensions of the existing methods would treat outliers identically to single object clusters.

The complement of some similarity measures (that is, $1 - s$ for a similarity measure s) for partitions are metrics, and it has been argued that a metric structure better matches our intuitions about similarity [22]. However, in practice, many of the most popular similarity measures, such as the ARI [14], NMI [17] (using the arithmetic mean, the default in the popular python machine learning library scikit-learn), and AMI [27], in addition to the Omega [5], oNMI [18], and ECS [10] from Section 2, are not the complement of a metric. Thus, even though being the complement of a metric is desirable, we conclude that this property is not absolutely necessary for a pragmatic similarity measure.

Proposition 2. *The complement of the proposed measure, $1 - F_{wo}^*$, is not a metric.*

Proof. We provide a counterexample for the triangle inequality. Let $M(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = 1 - F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)})$. Let $X = \{1, 2, 3\}$, and consider the following three clusterings:

$$\mathcal{C}^{(1)} = \{\{1, 2, 3\}\}, \quad \mathcal{C}^{(2)} = \{\{1\}, \{1, 2, 3\}\}, \quad \mathcal{C}^{(3)} = \{\{1\}, \{2, 3\}\}$$

We compute

$$\begin{aligned} M(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) &= 1 - \left(0.5(1) + 0.5 \left(\frac{1}{4} \frac{1}{3} + \frac{3}{4} \frac{1}{1} \right) \right) = \frac{1}{12}, \\ M(\mathcal{C}^{(2)}, \mathcal{C}^{(3)}) &= 1 - \left(0.5 \left(\frac{1}{4} \frac{1}{1} + \frac{3}{4} \frac{2}{3} \right) + 0.5 \left(\frac{1}{3} \frac{1}{1} + \frac{2}{3} \frac{2}{3} \right) \right) = \frac{17}{72}, \\ M(\mathcal{C}^{(1)}, \mathcal{C}^{(3)}) &= 1 - \left(0.5 \left(\frac{2}{3} \right) + 0.5 \left(\frac{1}{3} \frac{1}{3} + \frac{2}{3} \frac{2}{3} \right) \right) = \frac{7}{18}, \end{aligned}$$

and find that

$$M(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) + M(\mathcal{C}^{(2)}, \mathcal{C}^{(3)}) = \frac{1}{12} + \frac{17}{72} = \frac{23}{72} < \frac{7}{18} = M(\mathcal{C}^{(1)}, \mathcal{C}^{(3)}),$$

and the triangle inequality does not hold. \square

Remark. Since none of the clusterings used in the proof of Proposition 2 contain outliers, this counterexample shows that $1 - F_w^*$ is also not a metric.

Finally, we examine the effect of a small change to one of the clusterings being compared.

Theorem 3. *Let $\mathcal{C}^{(1)}$ and $\mathcal{C}^{(2)}$ be clusterings of X , and let $\tilde{\mathcal{C}}^{(1)}$ be a clustering created from $\mathcal{C}^{(1)}$ by moving one object, x , into or out of one cluster. Let $B = \sum_{C \in \mathcal{C}^{(2)}} |C|$. If $|\mathcal{C}^{(1)}| = |\tilde{\mathcal{C}}^{(1)}|$ (no clusters were created or destroyed), say x was added or removed from $C_1 \in \mathcal{C}^{(1)}$, then*

$$\begin{aligned} & \left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\ &= \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \gamma}{B} \right\} \right), \end{aligned}$$

where $\gamma = |\{C \in \mathcal{C}^{(2)} : C \text{ matched with } C_1 \text{ or } x \in C\}|$. Otherwise,

$$\left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| = \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \Gamma}{B} \right\} \right),$$

where $\Gamma = |\{C \in \mathcal{C}^{(2)} : x \in C\}|$.

The proof can be found in Appendix A. With Theorem 3, we confirm that, subject to some mild constraints on the clusterings, moving one element into or out of a cluster has only a small effect on the similarity to a reference clustering. In both cases, the right-hand side can be improved to $\mathcal{O}(\frac{1}{|X|})$ with additional mild constraints on $\mathcal{C}^{(2)}$ that prevent an object from being in too many clusters. Furthermore, in the proof of Theorem 3, we use the following Lemma that is of independent interest.

Lemma 4. *Let $A = \sum_{C \in \mathcal{C}^{(1)}} |C|$.*

$$\left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| = \mathcal{O} \left(\frac{1}{A} \right).$$

Remark. Lemma 4 implies that if A is sufficiently large, then $\hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)})$ is also robust to a small change in $\mathcal{C}^{(1)}$. Thus, this is a good choice for evaluating exploratory data analysis algorithms when comparing to a ground truth without outliers.

5 Experiments

5.1 Intuitive Properties

In this section, we recreate and extend an experiment from Gates et al. [10] to check for common biases in the behaviour of F_{wo}^* . The original experiment consists of three series of clustering comparisons that have an expected intuitive behaviour. None of the scenarios have outliers nor overlapping clusters, but a good measure should still match our intuitions for partitions.

For the first experiment, a clustering is created from 1024 objects with 32 clusters of each of size 32. The second clustering starts as a copy of the first, but a fraction of its objects have their cluster shuffled. We expect that as the fraction shuffled increases, the similarity should be monotone decreasing to some non-zero value.

In the second scenario, the first clustering is again 1024 objects split into 32 clusters of size 32. The second clustering starts with as a different clustering (randomly selected) with 32 clusters of size 32, and evolves by reassigning a random object to a new cluster chosen proportional to the cluster sizes. As the process proceeds, we expect the cluster sizes to become more heterogeneous (a few large clusters, many small clusters) and we again expect the similarity to be monotone decreasing to some non-zero value.

Finally, we test for bias in the number of clusters. The first clustering has 1024 objects split into 8 clusters of size 128, and the second clustering is chosen uniformly from partitions of 1024 objects into k equally sized clusters. As k increases, we expect the similarity to be monotone decreasing, and for the similarity to be almost 0 in the extreme case of 1024 clusters of size 1.

As seen in Figure 1, the behaviour of F_{wo}^* matches our intuition in each experiment.

Next, we test two more scenarios that involve clusterings with either outliers or overlapping clusters. To generate clusterings, we use a method from the recently proposed *Artificial Benchmark for Community Detection with Outliers and Overlaps* (ABCD+o²) [2]. This method leverages a low-dimensional geometric reference layer to create correlations between the overlapping clusters that are similar to those found in real world graph datasets.

As input, the method takes a set of initial cluster sizes p_1, p_2, \dots, p_m (such that for n objects, $\sum_{1 \leq i \leq m} p_i = n$) and an expansion factor η that corresponds to the average number of clusters per object. First, the method creates a partition P_1, P_2, \dots, P_m with the given sizes p_1, p_2, \dots, p_m . Each object is assigned an i.i.d. uniform random vector in the unit circle. Next, for each $1 \leq i \leq m$, the object with the vector furthest from the origin that has not already been assigned to a part become the seed of part i . The part P_i is set as the seed object and its $p_i - 1$ nearest neighbours that have also not yet been assigned to a part. In the second stage, each part is expanded around its center of mass until it reaches size $p_i \cdot \eta$, where η is an input parameter governing the average number of clusters to which an object belongs. For an exact description of the method, we refer to [2].

In the fourth scenario, we test each comparison measure on overlapping clusters generated with the method described above. For 25 samples, we generate an underlying geometric layer and initial partition, and then compare various values of η to a reference clustering where

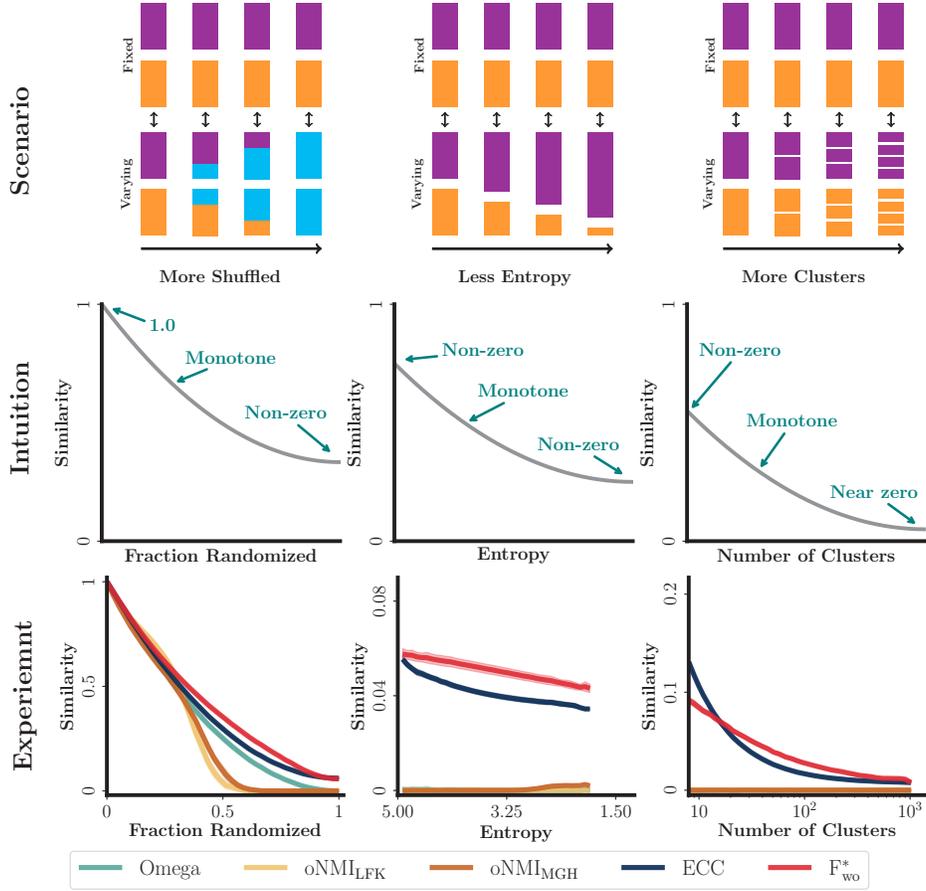


Figure 1: Recreation of the experiment in Figure 2 of [10] with F_{wo}^* . Each column represents one scenario, including a visual of the scenario, the intuitive behaviour of a similarity measure in this scenario, and the actual behaviour of several measures in each row respectively. Each line is the average of 100 simulations and the shaded region (usually too small to appear) covers plus or minus one standard deviation. In each scenario, the proposed F_{wo}^* measure matches our intuition.

$\eta = 3$. Since the second stage of the algorithm is deterministic, we expect perfect similarity when $\eta = 3$ and decreasing similarity as $|\eta - 3|$ increases.

In the fifth scenario, we reverse second stage of the method and shrink each cluster around its center of mass (this can be thought of as setting $\eta < 1$). Since we start with a partition, there will be outliers, but no overlapping clusters. We compare to a reference clustering corresponding to $\eta = 0.5$. This time we test two out of five measures, as only F_{wo}^* is explicitly defined for outliers, and ECS can accommodate outliers by treating them as isolated vertices in the cluster-induced element graph.

Once again, as seen in Figure 2, F_{wo}^* behaves as intuition would suggest. The ECS also follows the intuition in both scenarios, although the similarity for highly overlapping clusterings (large η) is rather high.

5.2 Graph Aware Clustering Comparison

Since overlapping clusters are often necessary to fully capture the communities in complex networks [16], we provide a short discussion about incorporating the structure of the graph when comparing graph clusterings. In particular, Poulin and Th  berge [24] propose several graph-aware extensions of popular partition measures by transforming a partition of the vertices to a binary edge classification (inter-part versus intra-part). In the case of overlapping clusters, we cannot use a binary edge classification, but we can create an edge clustering from a node clustering. For a graph G with vertices V and edges E , and a clustering \mathcal{C} of the vertices, we define the edge clustering induced by \mathcal{C} as

$$\mathcal{E}_G(\mathcal{C}) = \{\{uv \in E : u, v \in C\} \forall C \in \mathcal{C}\}.$$

We can also define the vertex clustering induced by an edge clustering, as is commonly used by edge clustering algorithms [1],

$$\mathcal{C}_G(\mathcal{E}) = \{\{v \in V : \exists uv \in C\} \forall C \in \mathcal{E}\}.$$

When evaluating the performance of a graph clustering algorithm, we can compare clusterings from both the vertex and edge perspectives. However, as could be the case when evaluating a linkage based edge clustering algorithm [1, 6], an edge cluster may span a vertex cluster while only covering a small fraction of the internal edges. To remedy this, we can define the closure of an edge clustering $\bar{\mathcal{E}}_G$ to add these internal edges back to the edge cluster, which can be computed as $\bar{\mathcal{E}}_G = \mathcal{E}_G(\mathcal{C}_G(\mathcal{E}))$. When using a graph-aware measure for comparing an edge clustering algorithm to a ground truth vertex clustering, we recommend using the closed edge clustering since the internal edges will be included in the edge clustering induced by the ground truth.

In Figure 3, we evaluate three resolutions of the Leiden [26] graph clustering algorithm on synthetic *Artificial Benchmark for Community Detection* (ABCD) [15] graphs in an experiment similar to that in Figure 5 of [24]. The ABCD model, similarly to the popular LFR model [19], features power-law degrees and community sizes, but has a noise parameter $\xi \in [0, 1]$ that controls the proportion of edges in the background graph and allows for a smooth transition from disjoint communities to a completely random graph. The different resolutions of Leiden produce different cluster granularities, with 3.75 giving a fine clustering (many small clusters), 0.5 producing a coarse clustering (few large clusters), and the default resolution of 1.0 producing clustering with medium-sized clusters. While we see the same ranking from both vertex and edge perspectives for low noise (up to $\xi \approx 0.5$), the rankings are inverted for very high ξ . In a very noisy graph, most of the edges are inter-cluster, and thus the ground truth edge clustering has many outliers. Any fine partition will also have mostly inter-cluster edges, so it is similar to the ground truth from an edge clustering perspective, even if the vertex clusterings are not similar. For evaluating graph clustering algorithms, we recommend analyzing performance from both perspectives to strengthen the conclusions.

6 Conclusion

In this note, we defined F_{wo}^* , a general method for quantifying the similarity between two clusterings of a common set of objects when the clusterings may have outliers or overlapping clusters. The score has intuitive behaviour, is easy to compute, and does not exhibit some common biases, which makes it viable for evaluating clustering algorithms. In Appendix B, we provide a simple yet efficient python implementation of F_{wo}^* .

Code Availability

A python implementation with support for additional input types and graph aware comparisons is available at <https://github.com/ryandewolfe33/Fstar>. The repository also includes code used in the experiments of Section 5.

Acknowledgements

R.D. acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) via a CGS-M scholarship.

References

- [1] Yong-Yeol Ahn, James P. Bagrow, and Sune Lehmann. Link communities reveal multi-scale complexity in networks. *Nature*, 466:761–764, 2010. doi:10.1038/nature09182.
- [2] Jordan Barrett, Ryan DeWolfe, Bogumił Kamiński, Paweł Prałat, Aaron Smith, and François Théberge. The artificial benchmark for community detection with outliers and overlapping communities (ABCD+o²). In *Modelling and Mining Networks*, volume 15669 of *WAW 2025. Lecture Notes in Computer Science*, pages 125–140, Vilnius, Lithuania, 2025. doi:10.1007/978-3-031-92898-7_9.
- [3] Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining*, volume 7819 of *Lecture Notes in Computer Science*, pages 160–172. Springer, Berlin, Heidelberg, 2013. doi:978-3-642-37456-2_14.
- [4] Peter Christen, David J. Hand, and Nishadi Kirielle. A review of the F-measure: Its history, properties, criticism, and alternatives. *ACM Computing Surveys*, 56(3), 2023. doi:10.1145/3606367.
- [5] Linda M. Collins and Clyde W. Dent. Omega: A general formulation of the rand index of cluster recovery suitable for non-disjoint solutions. *Multivariate Behavioral Research*, 23(2):231–242, 1988. doi:10.1207/s15327906mbr2302_6.

- [6] Ryan DeWolfe. Hierarchical single-linkage clustering for community detection with overlaps and outliers, 2025. doi:10.48550/arXiv.2509.02334.
- [7] Ryan DeWolfe and Jeffrey L. Andrews. Random models for adjusting fuzzy rand index extensions. *Advances in Data Analysis and Classification*, 19:361–385, 2025. doi:10.1007/s11634-025-00625-w.
- [8] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010. doi:10.1016/j.physrep.2009.11.002.
- [9] Alexander J. Gates and Yong-Yeol Ahn. The impact of random models on clustering similarity. *Journal of Machine Learning Research*, 18(87):1–28, 2017. URL: <http://jmlr.org/papers/v18/17-039.html>.
- [10] Alexander J. Gates, Ian B. Wood, William P. Hetrick, and Yong-Yeol Ahn. Element-centric clustering comparison unifies overlaps and hierarchy. *Scientific Reports*, 9(1):8574, 2019. doi:10.1038/s41598-019-44892-y.
- [11] Martijn M Gösgens, Alexey Tikhonov, and Liudmila Prokhorenkova. Systematic analysis of cluster similarity indices: How to validate validation measures. In *International Conference on Machine Learning*, volume 139 of *PMLR*, pages 3799–3808, 2021. URL: <https://proceedings.mlr.press/v139/gosgens21a.html>.
- [12] David J. Hand, Peter Christen, and Nishadi Kirielle. F*: an interpretable transformation of the F-measure. *Machine Learning*, 110(3):451–456, 2021. doi:10.1007/s10994-021-05964-1.
- [13] C. Hennig, M. Meila, F. Murtagh, and R. Rocci. *Handbook of Cluster Analysis*. Chapman and Hall/CRC, New York, 2015. doi:10.1201/b19706.
- [14] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218, 1985. doi:10.1007/BF01908075.
- [15] Bogumił Kamiński, Paweł Prałat, and François Théberge. Artificial benchmark for community detection (ABCD)—fast random graph model with community structure. *Network Science*, pages 1–26, 2021. doi:10.1017/nws.2020.45.
- [16] Bogumil Kaminski, Pawel Prałat, and François Théberge. *Mining complex networks*. Chapman and Hall/CRC, 2021. doi:10.1201/9781003218869.
- [17] Tarald O. Kvalseth. Entropy and correlation: Some comments. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(3):517–519, 1987. doi:10.1109/TSMC.1987.4309069.
- [18] Andrea Lancichinetti, Santo Fortunato, and János Kertész. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11:033015, 2009. doi:10.1088/1367-2630/11/3/033015.

- [19] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. Benchmark graphs for testing community detection algorithms. *Physical review E*, 78(4):046110, 2008. doi: 10.1103/PhysRevE.78.046110.
- [20] Andrea Lancichinetti, Filippo Radicchi, José J. Ramasco, and Santo Fortunato. Finding statistically significant communities in networks. *PLoS ONE*, 6(4):e18961, 2011. doi: 10.1371/journal.pone.0018961.
- [21] Aaron F. McDaid, Derek Greene, and Neil Hurley. Normalized mutual information to evaluate overlapping community finding algorithms, 2013. doi:10.48550/arXiv.1110.2515.
- [22] Marina Meilă. Comparing clusterings—an information based distance. *Journal of Multivariate Analysis*, 98(5):873–895, 2007. doi:10.1016/j.jmva.2006.11.013.
- [23] Minhyuk Park, Yasamin Tabatabaee, Vikram Ramavarapu, Baqiao Liu, Vidya Kamath Pailodi, Rajiv Ramachandran, Dmitriy Korobskiy, Fabio Ayres, George Chacko, and Tandy Warnow. Identifying well-connected communities in real-world and synthetic networks. In *Complex Networks & Their Applications XII*, volume 1142 of *Studies in Computational Intelligence*, pages 3–14. Springer, Cham, 2024. doi:10.1007/978-3-031-53499-7_1.
- [24] Valérie Poulin and François Théberge. Comparing graph clusterings: Set partition measures vs. graph-aware measures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(6):2127–2132, 2021. doi:10.1109/TPAMI.2020.3009862.
- [25] Mohammad Rezaei and Pasi Fränti. Set matching measures for external cluster validity. *IEEE Transactions on Knowledge and Data Engineering*, 28(8):2173–2186, 2016. doi: 10.1109/TKDE.2016.2551240.
- [26] Vincent A. Traag, Ludo Waltman, and Nees Jan van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1), 2019. doi:10.1038/s41598-019-41695-z.
- [27] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, 11(95):2837–2854, 2010. URL: <http://jmlr.org/papers/v11/vinh10a.html>.
- [28] Jierui Xie, Stephen Kelley, and Boleslaw K. Szymanski. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys*, 45(4):1–35, 2013. doi:10.1145/2501654.2501657.

A Proof of Theorem 3

Theorem (3). Let $\mathcal{C}^{(1)}$ and $\mathcal{C}^{(2)}$ be clusterings of X , and let $\tilde{\mathcal{C}}^{(1)}$ be a clustering created from $\mathcal{C}^{(1)}$ by moving one object, x , into or out of one cluster. Let $B = \sum_{C \in \mathcal{C}^{(2)}} |C|$. If $|\mathcal{C}^{(1)}| = |\tilde{\mathcal{C}}^{(1)}|$ (no clusters were created or destroyed), say x was added or removed from $C_1 \in \mathcal{C}^{(1)}$. Let $\gamma = |\{C \in \mathcal{C}^{(2)} : C \text{ matched with } C_1 \text{ or } x \in C\}|$. Then

$$\begin{aligned} & \left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\ &= \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \gamma}{B} \right\} \right). \end{aligned}$$

Otherwise, let $\Gamma = |\{C \in \mathcal{C}^{(2)} : x \in C\}|$, and

$$\left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| = \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \Gamma}{B} \right\} \right).$$

First, we prove a few lemmas about some of the intermediate scores, and we will then combine these results to prove Theorem 3. For rest of the section, let $\mathcal{C}^{(1)}$ and $\mathcal{C}^{(2)}$ be clusterings of X . Let $x \in X \setminus C_1$ for some $C_1 \in \mathcal{C}^{(1)}$, and create a new cluster $\tilde{C} = C_1 \cup \{x\}$. Let $\tilde{\mathcal{C}}^{(1)}$ be the clustering obtained from $\mathcal{C}^{(1)}$ by replacing C_1 with \tilde{C} . This corresponds to the first scenario in Theorem 3.

Lemma 5. If $C_j \subseteq X$ then

$$\left| F^*(C_1, C_j) - F^*(\tilde{C}, C_j) \right| \leq \frac{1}{|C_1|} \quad \text{and} \quad \left| F^*(C_1, C_j) - F^*(\tilde{C}, C_j) \right| \leq \frac{1}{|C_j|}.$$

Proof. There are two cases to consider. First, if $x \in C_j$, then

$$\left| F^*(C_1, C_j) - F^*(\tilde{C}, C_j) \right| = \left| \frac{|C_1 \cap C_j|}{|C_1 \cup C_j|} - \frac{|\tilde{C} \cap C_j|}{|\tilde{C} \cup C_j|} \right| = \frac{1}{|C_1 \cup C_j|}.$$

And if $x \notin C_j$, then

$$\left| F^*(C_1, C_j) - F^*(\tilde{C}, C_j) \right| = \left| \frac{|C_1 \cap C_j|}{|C_1 \cup C_j|} - \frac{|\tilde{C} \cap C_j|}{|\tilde{C} \cup C_j|} \right| = \frac{|C_1 \cap C_j|}{|C_1 \cup C_j| \cdot |\tilde{C} \cup C_j|} \leq \frac{1}{|\tilde{C} \cup C_j|}.$$

In either case, $|C_1|, |C_j| \leq |C_1 \cup C_j| \leq |\tilde{C} \cup C_j|$ so $1/|\tilde{C} \cup C_j| \leq 1/|C_1 \cup C_j| \leq 1/|C_1|, 1/|C_j|$. \square

Lemma (4). Let $A = \sum_{C \in \mathcal{C}^{(1)}} |C|$.

$$\left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| = \mathcal{O} \left(\frac{1}{A} \right).$$

Proof. Recall that

$$\hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) = \frac{1}{A} \sum_{C \in \mathcal{C}^{(1)}} |C| F^*(C, \mathcal{C}^{(2)}).$$

Using the relationship between $\mathcal{C}^{(1)}$ and $\tilde{\mathcal{C}}^{(1)}$, we get

$$\hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) = \frac{1}{A+1} \left((|C_1| + 1) F^*(C_1, \mathcal{C}^{(2)}) + \sum_{C_i \in \mathcal{C}^{(1)}, i \neq 1} |C_i| F^*(C_i, \mathcal{C}^{(2)}) \right).$$

It follows that

$$\begin{aligned} & \left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\ & \leq \left| \frac{|C_1| F^*(C_1, \mathcal{C}^{(2)})}{A} - \frac{(|C_1| + 1) F^*(\tilde{C}_1, \mathcal{C}^{(2)})}{A+1} \right| \\ & \quad + \left| \frac{1}{A} - \frac{1}{A+1} \right| \cdot \left| \sum_{C_i \in \mathcal{C}^{(1)}, i \neq 1} |C_i| F^*(C_i, \mathcal{C}^{(2)}) \right| \\ & \leq \frac{A|C_1| \left| F^*(C_1, \mathcal{C}^{(2)}) - F^*(\tilde{C}_1, \mathcal{C}^{(2)}) \right| + \left| |C_1| F^*(C_1, \mathcal{C}^{(2)}) - A F^*(\tilde{C}_1, \mathcal{C}^{(2)}) \right|}{A \cdot (A+1)} \\ & \quad + \frac{1}{A \cdot (A+1)} \cdot A. \end{aligned}$$

Using Lemma 5, we get

$$\left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \leq \frac{A + A + A}{A \cdot (A+1)} = \mathcal{O}\left(\frac{1}{A}\right),$$

which finishes the proof. \square

Lemma 6. Let $B = \sum_{C \in \mathcal{C}^{(2)}} |C|$ and $\gamma = |\{C \in \mathcal{C}^{(2)} : C \text{ matched with } C_1 \text{ or } x \in C\}|$.

$$\left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \tilde{\mathcal{C}}^{(1)}) \right| \leq \frac{\gamma}{B}.$$

Proof. First, note that if $x \notin C$, then $F^*(C, \tilde{C}) < F^*(C, C_1)$. Thus, any cluster not containing x that does not match with C_1 will not match with \tilde{C} . So, if $x \notin C$ and C is not matched with C_1 then $F^*(C, \mathcal{C}^{(1)}) = F^*(C, \tilde{\mathcal{C}}^{(1)})$. Thus, using Lemma 5 we get

$$\begin{aligned} \left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \tilde{\mathcal{C}}^{(1)}) \right| &= \left| \sum_{C \in \mathcal{C}^{(2)}} \frac{|C|}{B} \left(F^*(C, \mathcal{C}^{(1)}) - F^*(C, \tilde{\mathcal{C}}^{(1)}) \right) \right| \\ &\leq \sum_{C \in \mathcal{C}^{(2)}} \frac{|C|}{B} \left| F^*(C, \mathcal{C}^{(1)}) - F^*(C, \tilde{\mathcal{C}}^{(1)}) \right| \\ &\leq \frac{\gamma}{B}, \end{aligned}$$

and the proof is finished. \square

Now, we can prove Theorem 3.

Proof of Theorem 3. Due to symmetry, without loss of generality, we may assume one of the following four cases.

- 1a. $x \notin O^{(1)}$ and is added to existing cluster C_1 to make clustering $\tilde{\mathcal{C}}^{(1)}$.
- 1b. $x \in O^{(1)}$ and is added to existing cluster C_1 to make clustering $\tilde{\mathcal{C}}^{(1)}$.
- 2a. $x \notin O^{(1)}$ and makes a new cluster $\{x\}$ to make clustering $\hat{\mathcal{C}}^{(1)}$.
- 2b. $x \in O^{(1)}$ and makes a new cluster $\{x\}$ to make clustering $\hat{\mathcal{C}}^{(1)}$.

Recall that $A = \sum_{C \in \mathcal{C}^{(1)}} |C|$ and $B = \sum_{C \in \mathcal{C}^{(2)}} |C|$.

Case 1a:

$$\begin{aligned}
& \left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
&= \left| 0.5 \left(\frac{|X \setminus O^{(1)}|}{|X|} \right) \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) + 0.5 \left(\frac{|X \setminus O^{(2)}|}{|X|} \right) \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) \right. \\
&\quad \left. - 0.5 \left(\frac{|X \setminus O^{(1)}|}{|X|} \right) \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) - 0.5 \left(\frac{|X \setminus O^{(2)}|}{|X|} \right) \hat{F}_w^*(\mathcal{C}^{(2)}, \tilde{\mathcal{C}}^{(1)}) \right| \\
&\leq \frac{|X \setminus O^{(1)}|}{|X|} \frac{1}{A} + \frac{|X \setminus O^{(2)}|}{|X|} \frac{\gamma}{B},
\end{aligned}$$

by Lemmas 4 and 6. And, since $|X \setminus O^{(1)}| \leq A$, we conclude that

$$\left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \leq \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \gamma}{B} \right\} \right).$$

Case 1b: First, we consider each direction of \hat{F}_{wo}^* individually.

$$\begin{aligned}
& \left| \hat{F}_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
&= \left| \frac{|O^{(1)}|}{|X|} F^*(O^{(1)}, O^{(2)}) + \frac{|X \setminus O^{(1)}|}{|X|} \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) \right. \\
&\quad \left. - \frac{|\tilde{O}^{(1)}|}{|X|} F^*(\tilde{O}^{(1)}, O^{(2)}) - \frac{|X \setminus \tilde{O}^{(1)}|}{|X|} \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
&\leq \frac{1}{|X|} F^*(O^{(1)}, O^{(2)}) + \frac{|O^{(1)}| - 1}{|X|} \left| F^*(O^{(1)}, O^{(2)}) - F^*(\tilde{O}^{(1)}, O^{(2)}) \right| \\
&\quad + \frac{1}{|X|} \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) + \frac{|X \setminus O^{(1)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
&= \mathcal{O} \left(\frac{1}{|X|} \right).
\end{aligned}$$

And for the other direction,

$$\begin{aligned}
& \left| \hat{F}_{wo}^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_{wo}^*(\mathcal{C}^{(2)}, \tilde{\mathcal{C}}^{(1)}) \right| \\
& \leq \frac{|O^{(2)}|}{|X|} \left| F^*(O^{(2)}, O^{(1)}) - F^*(O^{(2)}, \tilde{O}^{(1)}) \right| + \frac{|X \setminus O^{(2)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \tilde{\mathcal{C}}^{(1)}) \right| \\
& = \mathcal{O} \left(\frac{1}{|X|} \frac{|X \setminus O^{(2)}| \gamma}{B} \right).
\end{aligned}$$

So together,

$$\left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\tilde{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| = \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \gamma}{B} \right\} \right).$$

Preliminaries for Case 2a and 2b: For these cases, we need to slightly alter the proofs of Lemmas 4 and 6 for when x creates a new cluster. First, a similar bound to Lemma 4:

$$\begin{aligned}
& \left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
& = \left| \sum_{C \in \mathcal{C}^{(1)}} \frac{|C|}{A} F^*(C, \mathcal{C}^{(2)}) - \sum_{C \in \mathcal{C}^{(1)}} \frac{|C|}{A+1} F^*(C, \mathcal{C}^{(2)}) - \frac{1}{A+1} F^*(\{x\}, \mathcal{C}^{(2)}) \right| \\
& \leq \left(\sum_{C \in \mathcal{C}^{(1)}} \frac{1}{A(A+1)} + \frac{1}{A+1} F^*(\{x\}, \mathcal{C}^{(2)}) \right) \\
& = \mathcal{O} \left(\frac{1}{A} \right).
\end{aligned}$$

And also a similar bound to Lemma 6. In this case, if $x \notin C$, then $F^*(C, \{x\}) = 0$ so C will not match with $\{x\}$. Let $\Gamma = |\{C \in \mathcal{C}^{(2)} : x \in C\}|$ Thus,

$$\begin{aligned}
& \left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \hat{\mathcal{C}}^{(1)}) \right| \\
& \leq \sum_{C \in \mathcal{C}^{(2)}} \frac{|C|}{B} \left| F^*(C, \mathcal{C}^{(1)}) - F^*(C, \hat{\mathcal{C}}^{(1)}) \right| \\
& \leq \frac{\Gamma}{|B|}.
\end{aligned}$$

Case 2a:

$$\begin{aligned}
& \left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
& \leq \frac{|X \setminus O^{(1)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| + \frac{|X \setminus O^{(2)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \hat{\mathcal{C}}^{(1)}) \right| \\
& = \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \Gamma}{B} \right\} \right).
\end{aligned}$$

Case 2b:

$$\begin{aligned}
& \left| F_{wo}^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - F_{wo}^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
& \leq \frac{1}{|X|} F^*(O^{(1)}, O^{(2)}) + \frac{|O^{(1)}| - 1}{|X|} \left| F^*(O^{(1)}, O^{(2)}) - F^*(\hat{O}^{(1)}, O^{(2)}) \right| \\
& \quad + \frac{1}{|X|} \hat{F}_w^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) + \frac{|X \setminus O^{(1)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(1)}, \mathcal{C}^{(2)}) - \hat{F}_w^*(\hat{\mathcal{C}}^{(1)}, \mathcal{C}^{(2)}) \right| \\
& \quad + \frac{|O^{(2)}|}{|X|} \left| F^*(O^{(2)}, O^{(1)}) - F^*(O^{(2)}, \hat{O}^{(1)}) \right| + \frac{|X \setminus O^{(2)}|}{|X|} \left| \hat{F}_w^*(\mathcal{C}^{(2)}, \mathcal{C}^{(1)}) - \hat{F}_w^*(\mathcal{C}^{(2)}, \hat{\mathcal{C}}^{(1)}) \right| \\
& = \mathcal{O} \left(\frac{1}{|X|} \max \left\{ 1, \frac{|X \setminus O^{(2)}| \Gamma}{B} \right\} \right).
\end{aligned}$$

□

B Python Implementation

```
import numpy as np
import scipy.sparse as sp

def compare(c1, c2):
    # Inputs c1, c2: The clusterings to compare.
    # Both must be scipy csr matrices of size (n_clusters, n_objects)
    # with c[i,j] = 1 if object j is in cluster i.
    if c1.shape[0] == 0 and c2.shape[0] == 0: # No clusters
        return 1
    c1_sizes = np.asarray(c1.sum(axis=1)).reshape(-1)
    c1_props = c1_sizes / np.sum(c1_sizes)
    c2_sizes = np.asarray(c2.sum(axis=1)).reshape(-1)
    c2_props = c2_sizes / np.sum(c1_sizes)
    intersect = c1.astype("float64") @ c2.transpose().astype("float64")
    row_sizes = sp.diags(c1_sizes)
    col_sizes = sp.diags(c2_sizes)
    nz = intersect.copy() # Store non-zero data indices
    nz.data[:] = 1
    union = row_sizes @ nz + nz @ col_sizes - intersect
    union.data = 1 / union.data
    fs = intersect.multiply(union)
    fs_l = fs.max(axis=1).toarray().reshape(-1)
    fs_lw = np.sum(fs_l * c1_props)
    fs_r = fs.max(axis=0).toarray().reshape(-1)
    fs_rw = np.sum(fs_r * c2_props)
    c1_o = c1.count_nonzero(axis=0) == 0
    c2_o = c2.count_nonzero(axis=0) == 0
    o_intersect = np.sum(c1_o * c2_o)
    o_fs = 0
    if o_intersect > 0:
        o_fs = o_intersect / (np.sum(c1_o) + np.sum(c2_o) - o_intersect)
    alpha = 1 - np.sum(c1_o) / len(c1_o)
    beta = 1 - np.sum(c2_o) / len(c2_o)
    fs_l_wo = (1 - alpha) * o_fs + alpha * fs_lw
    fs_r_wo = (1 - beta) * o_fs + beta * fs_rw
    return 0.5 * (fs_l_wo + fs_r_wo)
```

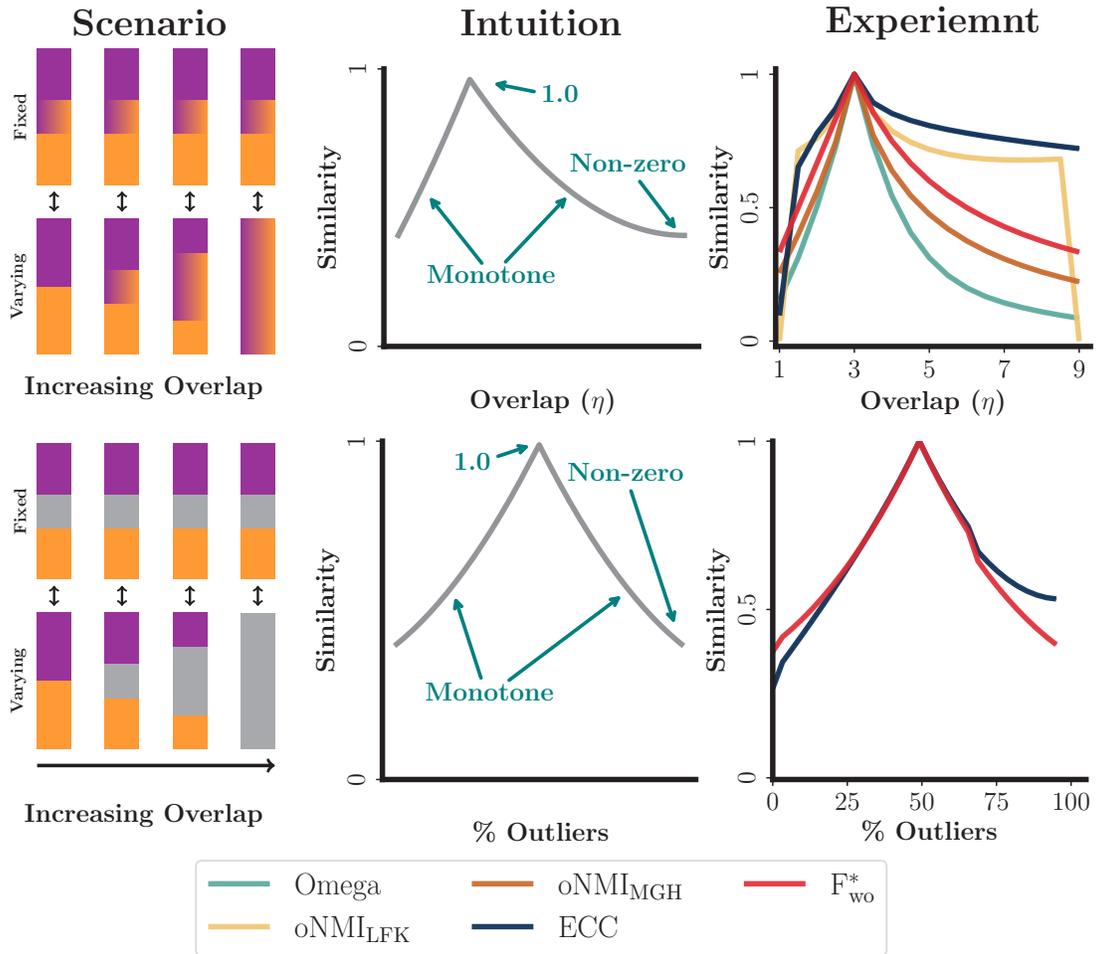


Figure 2: Comparing comparison measure behaviour on intuitive scenarios when overlaps or outliers are present. The top row is the scenario with overlapping clusters, and the second row has outliers; the columns, from left to right, show a visual of the scenario, the intuitive behaviour, and the actual behaviour of each method.

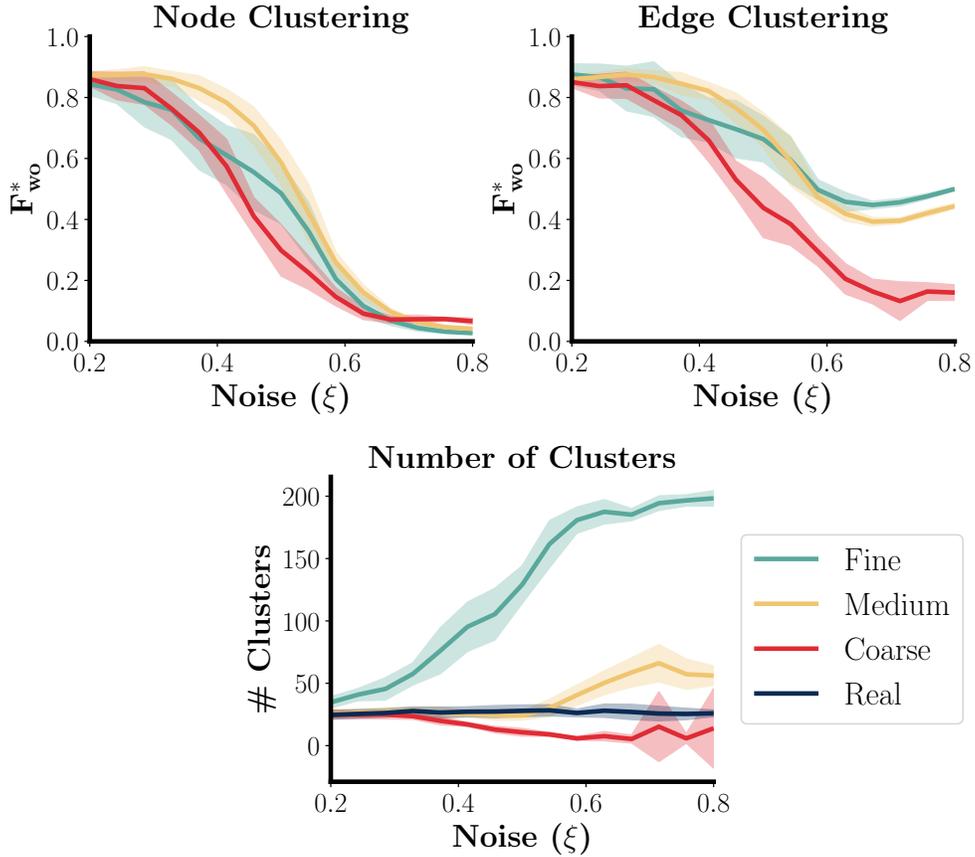


Figure 3: An experiment comparing graph aware clustering comparison, similar to that of Figure 5 from [24]. We use three resolutions (0.5, 1.0, and 3.75 shown in red, green, and yellow respectively) of the Leiden clustering algorithm on ABCD graphs with 2000 vertices and show, from left to right, the similarity of each detected vertex clustering to the ground truth vertex clusters, the similarity of the detected induced edge clustering to the ground truth induced edge clusters, and the number of clusters. The solid line represents the average over 5 samples and the shaded region covers plus or minus one standard deviation.