



Automatic Quasi-Clique Merger Algorithm – A hierarchical clustering based on subgraph-density[☆]



Scott Payne^a, Edgar Fuller^b, George Spirou^c, Cun-Quan Zhang^{a,*}

^a West Virginia University, United States of America

^b Florida International University, United States of America

^c University of South Florida, United States of America

ARTICLE INFO

Article history:

Received 5 June 2021

Received in revised form 20 August 2021

Available online 24 September 2021

Keywords:

Clustering

Agglomerative clustering

Community detection

Graph density

QCM

Facebook

Social network

ABSTRACT

The Automatic Quasi-Clique Merger algorithm is a new algorithm adapted from early work published under the name QCM (introduced by Ou and Zhang (2007)). The AQCM algorithm performs hierarchical clustering in any data set for which there is an associated similarity measure quantifying the similarity of any data i and data j . Importantly, the method exhibits two valuable performance properties: (1) the ability to automatically return either a larger or smaller number of clusters depending on the inherent properties of the data rather than on a parameter. (2) the ability to return a very large number of relatively small clusters automatically when such clusters are reasonably well defined in a data set. In this work we present the general idea of a quasi-clique agglomerative approach, provide the full details of the mathematical steps of the AQCM algorithm, and explain some of the motivation behind the new methodology. The main achievement of the new methodology is that the agglomerative process now unfolds adaptively according to the inherent structure unique to a given data set, and this happens without the time-costly parameter adjustment that drove the previous QCM algorithm. For this reason we call the new algorithm *automatic*. We provide a demonstration of the algorithm's performance at the task of community detection in a social media network of 22,900 nodes.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

The task of finding a “natural” or automatically determined number of clusters in an arbitrary data set has long been known impossible in terms of provable optimality [1]. That is, it is known that there cannot be one mathematical definition of optimality and one algorithm that finds the true optimal solution for all clustering tasks and data sets. Nevertheless, in both science and industrial data mining applications it is often necessary to discover, within a data set, clusters that have some measurable degree of quality in terms of their distinctness and the way they characterize the data according to some aspect prescribed by the application [2]. Many methods have been studied over decades and our work does not aim to present a survey so in our treatment of the background of this field we will mention only a few examples of the main approaches to automatic clustering.

[☆] This work was partially supported by an National Institute of Health (United States) Grant R01 DC015901 and an National Science Foundation (United States) Grant DMS-1700218.

* Corresponding author.

E-mail address: Cun-quan.zhang@mail.wvu.edu (C.-Q. Zhang).

1.1. AQCM

The automatic quasi-clique merger (AQCM) algorithm is a rework of the earlier quasi-clique merger (QCM) algorithm [3–6]. An explanation of AQCM subroutines and illustration of their implementation is given in Section 3. Here we begin with a basic description. AQCM is an agglomerative hierarchical clustering approach that takes as input an $n \times n$ similarity matrix S (or equivalently, a list of $\binom{n}{2}$ values representing the similarities for each pair of data points¹). In this work the definition of a similarity matrix S (or the $\binom{n}{2}$ values) is understood to be any function $S : X \times X \rightarrow \mathbb{R}^+$ where X is the data set and \mathbb{R}^+ is the non-negative real numbers, and $S[i, j] = S[j, i]$ for all i, j . Additional assumptions on the nature of similarity data input to AQCM are discussed in Section 1.3. For a detailed study of similarity measures in clustering see [7]. The objective of AQCM's algorithm design is to allow clusters to grow locally from well chosen pairs of data points (called seeds) known to be members of the same cluster. Seed selection and cluster growth adapt to local properties in the similarity data so that data points may be in more than one seed pair when appropriate and clusters may grow to a size suitable for preserving relative local similarity density. Hence, the algorithm supports multi-membership (sometimes referred to as fuzzy clustering) and clustering where cluster sizes vary widely (some small clusters and some large clusters).² The process repeats iteratively, each time treating the set of grown clusters as a smaller data set representing a higher level in the hierarchical scheme. The algorithm is finished running when there is only one cluster encompassing the data. While agglomerative methods are common in hierarchical clustering, the unique design of the subroutines in AQCM drive its adaptive properties and relatively efficient running time. These characteristics and the utility we have observed in our experimentation lend evidence that the algorithm can prove an important tool for obtaining automatic clustering in large and complex scientific data sets.

The AQCM method not only produces a generalized hierarchy tree (see Section 1.4.3 for definition), but also a clustering output automatically determined by an optimization method drawing on the properties of the hierarchy tree as a graph (network). The optimization strategy was originally presented in [6] where *density drop* was defined as the drop in local cluster density moving up successive levels along the branches of the hierarchy tree. In that work [6], Qi, et al. proposed that an optimal clustering can be defined as the “cut”³ across the tree featuring the most dramatic drop in cluster density. The advantage of the strategy is that tools from network science, such as max flow/min cut, provide optimization solutions for such a cut, and the resulting cut is not restricted to one level of the tree, thus allowing it to adapt to differences in local topological properties of the data along different branches of the tree. New in the work presented here is that we employ a recently proven mathematical optimization [8,9] for the selection of such a cut. The new method was designed to increase the flexibility to discover a potentially large number of clusters in a data set. We have found the method to perform as intended, and it appears to be robust in its ability to return either a small or large number depending on the inherent structure of the input data.

To speak to the intended use of the AQCM approach it is helpful to mention common uses of hierarchical clustering. In some hierarchical methods, the generated hierarchy tree is intended simply as a source for several “resolution levels” of clustering from which the user should choose one clustering output by some target optimization function. The term resolution refers to the granularity of clusters at a given hierarchical level, many small clusters being fine grain (i.e. high resolution) fewer larger clusters at an upper hierarchy being coarse grain (i.e. lower resolution). Basic divisive (top-down), binary hierarchical methods are often used as such, while the Louvain method of community detection [10] is a more complex example of that use. In contrast, other methods aim to yield a hierarchy tree which characterizes complex systems of topological hierarchies inherent in a data set. Such hierarchical use is often described as “multi-resolution”, where the purpose is to study the relationships between levels and the way they might relate to other metadata in a scientifically relevant way. The method of Jeub et al. [11] is an effective example of this type of use in network data analysis.

In the context of the strategies mentioned above we can suggest the role for the AQCM algorithm is similar to the latter example while having the benefit of an optimal clustering as in the former. Although one could potentially use the automatically chosen clustering within a larger industrial data mining pipeline, AQCM might be best used in exploratory scientific data mining where not much is known about a given data set but complex hidden patterns of clustering are thought to exist. By visually and/or computationally interpreting the hierarchy tree output by AQCM a user could discover various clustering patterns that may correlate with other known metadata. Such an investigative approach often plays an important role in biology for example as described in [2]. We have found AQCM to perform very well in this role in a number of data types and scientific investigations. In Section 5 we provide a demonstration of the algorithm's performance at the task of community detection in a social media network of 22,900 nodes. The algorithm's ability to respond automatically to inherent local properties in a data set while at the same time maintaining a relative short runtime make it a powerful tool in applications where data size and complexity both present large challenges.

¹ The software we have developed requires the full $n \times n$ similarity matrix for the purpose of faster matrix computations in various sections of the algorithm, however the mathematical procedures of the algorithm require simply the $\binom{n}{2}$ similarity values.

² The algorithm discovers clusters with such properties when those properties are inherent in the data, the algorithm does not force such properties onto the clustering output.

³ The graph theoretic term “cut” is defined in Section 1.4.

1.1.1. Contributions of AQCM

In this subsection, some novel contributions, special features are briefly summarized.

(1) Free of parameter tuning: there is no threshold in the algorithm for selection (this avoids time-consuming testing) (Section 3).

(2) A fine grain output (a large number of clusters) is produced (when naturally occurring in a data set) by each loop of iteration of AQCM (Section 3) and Automatic Community Selection (Section 4): (a selection based on edge-cuts with the minimum average weight).

(3) A multi-level output with hierarchical structure are produced by repeating loops of iterations (see Fig. 3 in Section 5), or simply by interpreting the hierarchy tree output by running AQCM once.

(4) The dendrogram constructed from AQCM is density monotonic (Section 3) in general (unless the data set in some extreme structure, see discussion in Section 1.3).

(5) Because of the density monotonicity in the hierarchical dendrogram, the automatic community selection based on density drop becomes feasible (Sections 3 and 4).

1.1.2. Software

Implementations of the main algorithms and post processing subroutines featured in this work are provided here [12]. We chose MATLAB as the implementation platform due to its optimized performance for matrix operations. Many of the mathematical procedures in the subroutines of AQCM can be efficiently expressed as matrix operations. We found that MATLAB's combination of runtime speed and ease of dynamic memory implementation made it an ideal choice for the scalable scientific work for which we envision that AQCM is best suited.

1.2. Automatic clustering, a short background

The naive approach for automatic determination of clusters would be to repeatedly run the k -means algorithm many times over many different choices of cluster number (k) and then choose the clustering output with the highest k -means quality function score. In many applications, however, this strategy is unfeasible for a myriad of reasons, the most obvious being time constraints and the fact that we must presume that the data is best represented as a set of n -dimensional points (many data sets are not represented as such). But moreover, the situation becomes untractable in cases where a data set is known to present some large but unknown number of clusters, rendering futile the job of choosing the range of k values.

A pivotal work in the development of automatic clustering techniques was the *affinity propagation algorithm* [13], or AP for short. The algorithm provides an excellent example of modern strategies that harness the power of Bayesian-like networks and allow a dynamic process to gravitate to an equilibrium. Specifically, AP relies on an adapted version of *belief propagation* [14] that Frey/Dueck refer to as “loopy belief”. AP is in a class of strategies that rely on the notion that a data cluster has a center, often referred to as a *medoid*. Also key to AP is that the input to the algorithm is not the data itself but rather a *similarity matrix* which must be an $n \times n$ real valued matrix where the ij entry scores how similar are data i and data j .

A different approach for the determination of clusters by medoids is to define cluster centers as data points with an unusually high number of “nearby points” as determined by some measure of local scale versus global scale distance. For data sets from metric spaces, the algorithm proposed in [15] (also see [16]) is able to recognize non-spherical clusters and therefore, is a significant improvement of k -means clustering. Additionally, the method in [15] allows for the number of clusters to emerge naturally, although user interaction and threshold choice is required.

Another class of clustering strategies relies only indirectly on the notion of a medoid and treats clustering as the problem of fitting a finite mixture model to the input data set. Many algorithms exist but a particularly effective method found here [17] builds upon various strategies that optimize log-likelihood and can be adapted to various statistical models (Gaussian being the default).

A final example category of clustering strategies is found in the graph theoretic task of *community detection*. Any similarity matrix representing a data set can be thought of as a weighted graph (network) where each vertex (node) represents a data point. Communities are groups of vertices within which the edge weights (similarities) are relatively high and edge weights between communities tend to be lower. So community detection is equivalent to clustering when an appropriate similarity matrix exists. There is a large literature on community detection. The main early example is found in the work of Girvan/Newman [18]. The methods proposed by Girvan/Newman are driven by the optimization of a target function called *modularity*. A key advancement in modularity optimization techniques was the algorithm called *Louvain method* [10] which allowed for very large data sets to be processed in feasible time on an appropriate system. A hierarchical technique that generalizes the Louvain method is found here [11] and a very recent modularity generalization algorithm is found here [19]. Importantly, that work (the latter) provides theorems proving that the quality of the clusters (according to a user chosen target function and resolution parameter) is optimized by further iteration.

As mentioned above, the review of methods provided here is only intended to illustrate some of the existing strategies. We have focused on strategies that aim to produce clusters where the number suits the unique properties of the input data set in some “natural” way without the requirement that the user have foreknowledge of the “correct” number of clusters to find. It is important to take a moment to note, however, that in many cases a data set may have multiple equally valid solutions to the clustering task or none (for discussions see [20] and [21]). Furthermore some clustering tasks, such as those designed to divide data for distributed computing, are such that the user *knows* the desired number of clusters ahead of time. Hence, in the next section we clarify the assumptions we make in this work regarding the nature of clustering tasks we would aim to perform and for which the algorithms discussed above are intended.

1.3. The clustering task

In this work we assume that a given clustering task on a data set has the following characteristics:

1. The data set has distinct structures in that there are subsets of the data within which the points tend to be more similar to each other and less similar to points in other subsets. The difficulty of such a clustering task would be determined by the range of similarities inside subsets and the range of similarities between subsets. The more that the between range overlaps the inside range, the more difficult it is to determine whether a pair of points are in the same cluster.
2. If the data are points in n -dimensional space we assume that the points are distributed such that there are some number of regions of higher point density separated by regions of lower density. The difficulty of such a clustering task would be determined by the distinctness of the dense regions against the space between regions and also the shape of the regions (which need not be symmetrical or spherical). For example the moon and sun problem and the target problem are construction methods that create data sets which will confound most modern clustering algorithms.

1.4. Notation and definitions

1.4.1. AQCM and graph theory

Many of the subroutines in the AQCM algorithm are adaptations or redesigns of subroutines from the QCM algorithm, and the general structure of each iteration of AQCM is the same as in QCM. Above we have described AQCM as a clustering algorithm that takes as input an $n \times n$ similarity matrix S , a common type of input structure in data mining [7] (see Section 1.1 for precise definition). It is important to note that the original publications of the QCM algorithm [3–6] presented the quasi-clique merger method as an algorithm for performing hierarchical “community detection”, a common data mining task in graph (network) data analysis. In that work, the input to QCM was a weighted graph (weighted network) and a community was defined as an “edge dense” set of vertices (nodes) where edge weights tend to be higher, edges between communities were expected to be more sparsely distributed and/or lower weight. The objective of the QCM algorithm and the type of input graph for which it was designed are equivalent to the idea of clustering in similarity data as commonly discussed in data mining today [7]. Thus far we have introduced the idea of AQCM using the language of similarity-based clustering due to its useful and accurate level of generality in framing the tasks for which AQCM may be used. But relevant to note is that the QCM algorithm itself is not able to operate on an unweighted graph, that is, the input must have weights associated with the edges which serve the same role as that of a similarity function. Hence, the description of the quasi-clique merger process in general might be most accurately described as a process that operates on similarity data.

On the other hand, while the mathematical steps in AQCM (and QCM) can be accurately described using the abstract notions of data points and similarity, many of the steps can actually be more easily discussed using their equivalent formulations in the language of graph theory. The graph theoretic concept of an *edge*, sometimes also referred to as a *link*, is a linguistically simple way to refer to the relationship between two data points, which in graph theoretic terms would be referred to as *vertices*, or sometimes *nodes*. Therefore, in the work presented here we may at times describe processes in terms of vertices and edges instead of data points and similarities, and input to AQCM may at times be described as a weighted graph instead of similarity values or a similarity matrix.

1.4.2. Graph theoretic terminology

For graph theoretic terminology and notation we follow [22–24] which are standard. A graph $G(V, E)$, or G for shorthand, is defined by a vertex set V (sometimes notated $V(G)$) and an edge set E (sometimes notated $E(G)$) where an edge $e \in E$ is a connection between some vertex $u \in V$ and some other vertex $v \in V$ so that e might also be notated uv . It is common to notate the size of the vertex set $|V| = n$ and the size of the edge set $|E| = m$. If G is a weighted graph then $w : E \rightarrow \mathbb{R}^+$ is the function assigning weights to edges and the weight of edge e is notated $w(e)$. Some authors assume that for a weighted graph all weights are non-zero or equivalently that $w(uv) = 0$ implies $uv \notin E$. In this work we will address such cases as necessary. If G is a directed graph (digraph) then each edge uv has the direction u to v where v is referred to as the *head* of the edge and u is the *tail*, and between two vertices u and v both the directed edges uv and vu might exist. Some authors use the term *arc* for a directed edge however here we will simply use the term edge and when the edge has a direction we will make that clear in context. Some additional technical graph-theoretic terminology is provided in [Appendix A.1](#).

As discussed in Section 1.4.1, parts of AQCM are sometimes more easily described using a graph theoretic language. Here we discuss notation. Let X be a data set and S be a similarity function $S : X \times X \rightarrow \mathbb{R}^+$ as defined in Section 1.1, then S is equivalent to a weighted graph $G(V, E)$ with weight function $w : E \rightarrow \mathbb{R}$ where vertex $v_i \in V$ is associated with data point $x_i \in X$ and for a pair of vertices v_i, v_j the edge weight $w(v_i v_j)$ is equal to the similarity value $S(x_i, x_j)$. Clearly the discrimination between $w(v_i v_j)$ and $S(x_i, x_j)$ is purely semantic, so henceforth, when we adopt a graph theoretic language of algorithm description we will use the notation S and refer to it as a weighted graph and refer to $S(x_i, x_j)$ as the edge weight for the edge $x_i x_j$.

1.4.3. Input and output structures of AQCM

The input to AQCM is an $n \times n$ similarity matrix S (or equivalently, a list of $\binom{n}{2}$ values representing the similarities for each pair of data points) where n is the number of data points in some data set X . The definition we adopt for the terminology *similarity* was given in Section 1.1 and further assumptions on the nature of similarity data input to AQCM are discussed in Section 1.3.

The output of a hierarchical clustering algorithm is typically represented by a dendrogram which is simply a diagram of a hierarchy tree. A precise graph theoretic definition of the term *tree* is given in Appendix A.1. A graph that is a tree is typically given the notation T instead of G . The *rooted tree* (see Appendix A.1 for precise definition) is the proper graph theoretic object that characterizes the concept of a hierarchy tree. The root represents the data as one set, the leaves represent the data as individuals, and nodes along a path from the root to a leaf represent cluster assignment for that leaf at various hierarchical levels. In a dendrogram, the depiction is typically oriented with the root at the top of the diagram and the leaves (representing individual data points) at the bottom. In data science, hierarchical relationships are often described using the terminology *parent* and *child*. In a hierarchy tree T , a parent node $v \in V(T)$ represents a cluster at some level (v is not a leaf node). For a node $u \in V(T)$ such that there is an edge $vu \in E(T)$ we say u is a child node of the parent node v . The node u represents a sub-cluster (of the parent cluster represented by v) at the next level down in the hierarchical scheme (or an individual data point if the parent is a level 1 cluster).

As mentioned in Section 1.1, the hierarchical information output by AQCM supports multimembership. That is, when appropriate, a data point (or cluster at some level in the hierarchy) could be a member of more than one cluster. In parent/child terminology we say a multimember child node has more than one parent node. Hence the hierarchy “tree” output by AQCM may not be a true tree according to the graph theoretic definition *tree* (see Appendix A.1). For simplicity of language in this work we will nonetheless refer to the hierarchical output produced by AQCM as a tree since we have clarified here the context and the idea of a generalized hierarchy tree (mentioned in Section 1.1) that allows for multimembership.

Based on the above we should mention that, in our method, the automatically chosen clustering selected from the hierarchy tree is not necessarily a *partition* of the data set (although it could be) but is more generally a *family* of subsets of the data notated $\mathcal{C} = \{C_1, \dots, C_k\}$ where $C_i \subset X$ is the i th cluster. Also, the clustering output \mathcal{C} may or may not form a *covering* of the data set as some outlier points might not be assigned to any cluster. We will refer to such unassigned points as *singletons* or simply *unclustered data*. Examples are seen in Section 2.

A final aspect of our method that we will clarify here is the relationship between automatic cluster selection and the graph theoretic notion of an optimal *edge cut*. The terminology *edge cut* is defined precisely in Appendix A.1, but here we can describe its relevance in a hierarchy tree as a rooted tree.⁴ Consider that every cluster C defined within the structure of hierarchy tree T is represented by a child node of some parent node (possibly the root). Hence, for any clustering output $\mathcal{C} = \{C_1, \dots, C_k\}$ selected from the tree T , there is an associated set of directed edges $H = \{e_1, \dots, e_k\} \subset E(T)$ where $e_i = v_i u_i$ is associated with cluster C_i through its representation by child node u_i . The set of edges H is called an *edge cut* because its removal separates the tree T into components. The selection of optimal edge cuts is a common topic in the theory of graphs and networks. Hence the selection of an optimal clustering $\mathcal{C} = \{C_1, \dots, C_k\}$ from hierarchy tree T may be approached by the selection of an optimal edge cut $H = \{e_1, \dots, e_k\} \subset E(T)$ provided that a suitable method exists for encoding, into the edges of T , information about the various clustering outputs available within T . We present our implementation of this approach in Section 4.

2. An illustrated introduction to AQCM

In this section we demonstrate the use of the AQCM algorithm to perform unsupervised cluster analysis on a basic abstract data set. The set was designed to exemplify the type of clustering tasks described in Section 1.3. 3D renderings of the data points are seen in Figs. 1-a,b,c. By mapping the data into the cube on the interval $[0,1]$, it can be interpreted as RGB color data (as illustrated in Fig. 1-a). The colors allow us to intuitively understand the inherent hierarchies and relationships between clusters.

2.1. The data set

The construction of the data set was as follows: we chose 9 3-dimensional points randomly and centered Gaussian distributions at those points. Covariances were chosen so that samples from the mixture model would generate point clouds with reasonably distinct boundaries as discussed in Section 1.3. Samples were drawn of random sizes from each Gaussian in the mixture model so as to generate planted clusters of different sizes (in terms of number of data points). There are 595 data points in total. The construction method is standard for synthetically generating point clouds with clustering properties similar to those often seen in experimentally collected scientific data.

⁴ For simplicity of explanation we describe the case of a tree with no multimembership, however the relationship exists similarly in trees featuring multimembership.

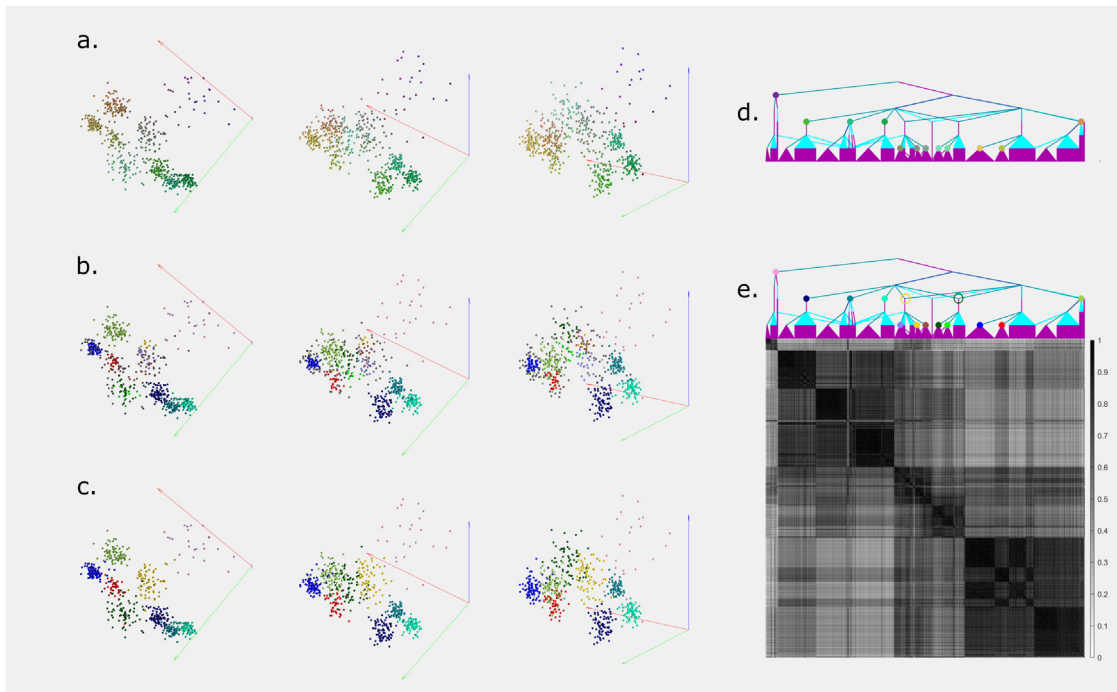


Fig. 1. **a.** 3-dimensional (RGB) data seen in three views, RGB axes are indicated by their color, each data point is colored by its RGB value, each axis range is $[0,1]$. **b.** Three views match Fig. 1-a, the data are colored to distinguish the clusters chosen by AQCM and match the indicator colors in the dendrogram seen in Fig. 1-e, data colored gray are “unclustered/outliers”. **c.** As in Fig. 1-b but here the colors distinguish the clusters obtained using a Gaussian finite mixture model fitting algorithm described in [17], we may compare and contrast this clustering with that obtained by AQCM. **d.** The dendrogram output by the AQCM algorithm, the layout of the dendrogram was computed using Matlab’s “layered layout” implementation for directed graphs, the colors of the lines indicate the similarity density drop between hierarchical levels according to the function described in Section 4 with light blue being more extreme dark blue moderate and purple least extreme, the colored dots indicate the automatically chosen clustering and their colors are the mean color per cluster. **e.** The dendrogram as described in Fig. 1-d except that here the cluster indicators are colored to match the illustration in Fig. 1-b, below the dendrogram the similarity data is displayed as organized by the dendrogram (this type of display is often referred to as a “heat map”) and the scale is given at the right, the yellow and green open circles show two alternate clusters represented in the AQCM output and these are essentially the same as the yellow and green clusters seen in Fig. 1-c (only differing by a few data).

2.2. Benchmark algorithm

For a benchmark/comparison we used the finite mixture model fitting algorithm of Figueiredo et al. described in [17] to automatically determine the cluster assignments illustrated by color in Fig. 1-c. The algorithm discovers 9 clusters whose centers and geometry match the underlying model used to generate the data, so we may think of this clustering as “correct” in that sense. We have tested this algorithm extensively in similar data sets and have found that it appears to return an accurate description of the planted clusters as long as the number of clusters is not too large and the differences in their relative geometries not too extreme. The algorithm is essentially parameter free for simple data sets with relatively few clusters since the required input range of cluster numbers to try can be fixed to an interval such as $[1,25]$ (that is what we used here). We have found that when the properties of the planted clusters become more extreme (in terms of number, size and geometry) the algorithm tends to produce less accurate results even if a good parameter range is used. We should also note that we have not benchmarked this algorithm’s performance in high dimensional data, so we cannot say whether it is feasible for such data in terms of accuracy or run time.

2.3. AQCM algorithm output

Illustrations of AQCM output are seen in Figs. 1-b,d,e. Figs. 1-b,c are color-coded diagrams of the clustering outputs of AQCM and the benchmark algorithm respectively.⁵

As AQCM requires a similarity matrix, we defined similarity as the negative of Euclidean distance shifted into the range 0 to 1. The similarity data is seen in Fig. 1-e as indexed by the hierarchy tree output by AQCM (seen above the matrix).

We begin by observing that the hierarchy tree output by AQCM has, on a high level, three main clusters (Fig. 1-d shows these): the purple dot indicates the average color of the cluster on the left branch while the right branch splits

⁵ We have “aligned” the colors in the illustration for ease of comparison between the two clustering outputs.

into two subclusters. Colored dots of further subclusters indicate that the two containing clusters represent a green group and a yellow/orange group with brown/tan/bluegrey clusters having multimembership of both green and yellow/orange. In the visual display of the color data seen in Fig. 1-a it is certainly clear that splitting the data in two between purple and non-purple data is a valid high level distinction and that the non-purple data can be described as yellow/orange and green with brown/tan/bluegrey in the middle. Hence the general structure of the hierarchy tree captures the inherent nature of the data well.

Using Figs. 1-b,c,e we can examine the clustering output selected by edge-cut cluster/community detection. The colors in Figs. 1-b,e illustrate that clustering. We have synchronized our labeling colors for Figs. 1-b,c so that the AQCM clustering may be directly compared with the benchmark algorithm clustering described above. The yellow/orange group and the green group each contain clusters that match the benchmark algorithm almost perfectly. The differences here are on account that AQCM leaves some data unclustered (see the red labeled and blue labeled clusters). This aspect of AQCM is a natural occurrence when a set of densely clustered points has outliers, AQCM may prefer the denser subset since it optimizes for cluster similarity density.

The main difference between the chosen clusters of edge-cut detection and the benchmark algorithm are in the treatment of the green labeled and yellow labeled clusters seen in Fig. 1-c. As illustrated by the open circles in Fig. 1-e, AQCM captures these same clusters at a higher hierarchy. Examination of the color data itself in Fig. 1-a would seem to indicate that the distinctions made by AQCM to split these two clusters further could also be considered valid since they each respectively contain locally dense similarity subsets as seen in Fig. 1-e. We should remark here that this difference between AQCM output and the benchmark algorithm demonstrates the different interpretations of data clusters given when one algorithm determines structure based on distance and another attempts to fit a Gaussian distribution. Hence the differences between the two outputs are inherent in the different definitions of the feature space of the data. Overall, the color data example demonstrates AQCM's ability to function well even within a feature space that presents challenges in terms of fuzzy clusters and regions of varying data density. One might expect that if the similarity function on the data were able to reflect the planted Gaussian clusters even more clearly then the clustering detected by the edge-cut algorithm might even more precisely match that of the benchmark algorithm. In general, the description of the data set provided by the hierarchy tree output does provide a robust multi-level interpretation of the planted structures as we might naturally see them with our human ability to perceive groups and dense point clouds.

3. AQCM algorithm and subroutines

In Section 3.1 we provide a basic description of each of the subroutines of AQCM so as to facilitate an easier understanding of their objectives and mathematical strategies. An illustration of the process is seen in Fig. 2. Full mathematical details are given in technical notation in Section 3.2. As discussed in Section 1.4.1, parts of AQCM are sometimes more easily described using a graph theoretic language. Here we adopt that style to describe the subroutines and their unfolding. See Section 1.4.2 for clarification of the relationship between our notation and graph theoretic language.

3.1. AQCM algorithm structure

AQCM has four subroutines: *seed selection*, *growth step*, *adjustment step*, *contraction step*. The AQCM algorithm iterates by running the subroutines in that order, each outputting objects for input to the next subroutine. The contraction step creates the input graph for the seed selection of the next iteration.

Input: $S : X \times X \rightarrow \mathbb{R}^+$ a similarity function on a data set X , equivalently, a weighted graph on a vertex set X where $S(x_i, x_j)$ is the edge weight between vertices x_i and x_j , in the first iteration of AQCM this is the current graph, though in the notation below we use the letter V instead of X in order to generalize the language for all iterations.

Output: T a generalized hierarchy tree, (see Section 1.4.3 for terminology)

Seed Selection: Select a subset $E_{seeds} \subset E$ (E is the edge set of the current graph on a vertex set V) such that for $e = v_i v_j \in E_{seeds}$, v_i and v_j are likely to be in the same cluster. This likelihood is ensured by estimations of locally close vertices made at each vertex: a seed edge $e = v_i v_j \in E_{seeds}$ is an edge such that in the estimations at vertex v_i and vertex v_j , each designated the other as a locally close vertex.

Growth Step: Grow a set of clusters $C = \{C_1, \dots, C_k\}$ by optimally adding vertices of V to individual seed edges of E_{seeds} . Growth is controlled locally in that for a given cluster C growing from a seed $e \in E_{seeds}$, C stops growing when its edge weight density falls below a threshold developed dynamically from the cluster's growth properties. The process unfolds linearly starting

with the highest weight seed edge in E_{seeds} . As a cluster C grows covering other seed edges, those edges are removed from E_{seeds} eliminating unnecessary processing.

Adjustment Step: Adjust the clustering $C = \{C_1, \dots, C_k\}$ when there are subsets of C with relatively large vertex overlap. Such subsets are possible since the seed/growth process may approximate a potential cluster several slightly different ways. In order to ensure that the most locally dense possible cluster is not overtaken by a less optimally dense approximation, the first iteration uses a different adjustment method than subsequent iterations. **First Iteration:** working through $\{C_1, \dots, C_k\}$ in order prioritizing whichever cluster C_i is most edge weight dense, remove any other cluster that overlaps C_i too much. Processing continues until degenerate overlaps are resolved. **Subsequent Iterations:** working through $\{C_1, \dots, C_k\}$ in order prioritizing whichever cluster C_i is most edgeweight dense, merge into C_i any other cluster that overlaps C_i too much. When more than one merge is possible, choose the merge resulting in the highest edgeweight density. Processing continues until degenerate overlaps are resolved. Note that for iterations two and higher, merge decisions are made by interpreting clusters as subsets of the data set X as opposed to subsets of contracted graph vertices (see Fig. 2 second row last column).

Contraction Step: Create a weighted graph representing the similarities between each pair of clusters C_i and C_j . As in the adjustment step, clusters are interpreted as subsets of the data set X as opposed to subsets of contracted graph vertices. Each cluster C_i is represented as a vertex v_i in the new graph, the edge weight between vertices v_i and v_j is the mean of all edgeweights of edges (of the graph on X) between C_i and C_j . The weighted graph created by the contraction step becomes the current graph which is the input to the next iteration of AQCM beginning with seed selection.

3.2. Subroutine details

In the following descriptions of the AQCM subroutines we use a more proper graph theoretic notation style. Let S be a weight function $S : E(K_n) \rightarrow \mathbb{R}^+$ where $E(K_n)$ is the edge set of a complete graph K_n on n vertices with vertex set V . We define the following.

Definition 3.1. For a subgraph Q of K_n induced by a vertex subset $C \subset V$, we define the density of Q by

$$\text{den}(Q) = \frac{2 \sum_{e \in E(Q)} S(e)}{|C|(|C| - 1)}$$

As Q is an induced subgraph of a complete graph, we have $|E(Q)| = \binom{|C|}{2}$. So it is easy to see that the function $\text{den}(Q)$ is simply the mean edgeweight over the subgraph Q , or equivalently, the average similarity between pairs of data in the cluster C .

Definition 3.2. For a vertex $v \notin V(Q)$, we define the contribution of v to Q by

$$\text{cont}(v, Q) = \frac{\sum_{u \in V(Q)} S(uv)}{|V(Q)|}.$$

The contribution $\text{cont}(v, Q)$ is the mean of edgeweights between v and Q , and as such it is clearly the simple and computationally efficient way to quantify a vertex's candidacy for joining a growing cluster.

3.2.1. Seed selection

Seed selection is a new subroutine in AQCM not found in published versions of QCM. Its purpose is to automatically select seeds from which to grow clusters in a way that adapts to the inherent properties of the input data. The intuition driving the seed selection algorithm comes from the assumptions we often make about what distinguishes a cluster within a data set (see Section 1.3). We expect that for a data x which is well within (not an outlier) a cluster C , the similarity values of x to other data in C would be in a numerical range that should be somewhat distinct from the numerical range of

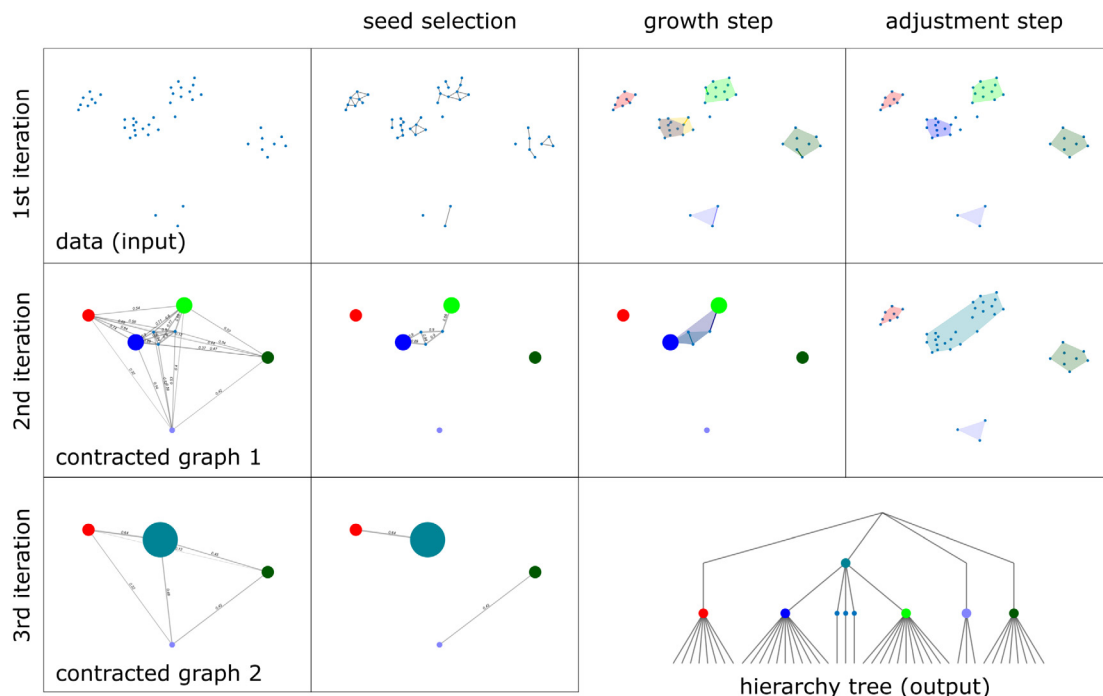


Fig. 2. An illustration of each subroutine’s output as AQC runs on a simple data set with planted clusters of 2D points (seen in the top left box). The similarity function we used was the “shifted negative of distance” as described in Section 2.3. Each row of the diagram shows an iteration (or hierarchical level) running through the processes of seed selection, growth, and adjustment. For each subroutine, its column in the diagram shows the relationship between the ways the algorithm views the data at different hierarchical levels. In column 1, iterations 2 and 3 show the *contracted graph* that will be input to seed selection, dot size reflects cluster size and their color is associated with steps in the previous iteration, the edge weights are seen to two decimal places. In illustrations of the growth step, the seeds from which the illustrated clusters grew are seen as colored edges. For this example data set, the 3rd iteration growth step resulted in one cluster covering the data, so the algorithm was finished at that step. Bottom right shows the dendrogram of the hierarchical clustering structure determined by the states of clustering at the end of each iteration, colored node indicators show the relations to the illustrated subroutines.

similarity values between x and the data not in C . In our seed selection algorithm we harness that property by considering, for data x , the list of similarity values to other data sorted in decreasing order. Thus we should expect some position in that ordered list where the similarity value abruptly drops as we move from similarities measured within C to those with other data outside C . While this property might not hold for every data x in C , it is easy to show that our algorithm will produce at least one seed edge inside each cluster and our algorithm will not choose edges representing a pair of data each in a different cluster assuming the similarity data has reasonable clusters as described in Section 1.3.

Seed selection algorithm:

Input: K_n a complete graph on n vertices, and edge weight function $S : E(K_n) \rightarrow \mathbb{R}^+$, when AQC is in the first iteration K_n and S are the inputs to AQC, in subsequent iterations K_n and S are objects created by the previous iteration contraction step.

Output: the seed edge set $E_{seeds} \subset E(K_n) = \{e_1, \dots, e_m\}$, for use in the growth step the set E_{seeds} should be output as an ordered set $\{e_{r_1}, \dots, e_{r_k}\}$ with $\{r_1, \dots, r_k\} \subset \{1, \dots, m\}$ and $S(e_{r_i}) \geq S(e_{r_{i+1}})$ for all i .

Step 1.

for each vertex $v \in V(K_n)$ **do**:

- order vertices decreasing by edgeweight with v :
- $\{u_1, \dots, u_{n-1}\}$ is the set of vertices not including v
- $\{u_{q_1}, \dots, u_{q_{n-1}}\}$ is an ordering with $S(vu_{q_i}) \geq S(vu_{q_{i+1}})$ for all i

compute a sequence of differences:

$$\{a_i\}_{i=1}^{n-2} \leftarrow \{S(vu_{q_i}) - S(vu_{q_{i+1}})\}_{i=1}^{n-2}$$

compute the median of that set (treated as a statistical sample):
 $M_a \leftarrow \text{median}(\{a_\ell \mid \ell \in \{1, \dots, n-2\}\})$

find the first significant edgewise drop:
 $t \leftarrow \min(\{i \mid a_i \geq M_a\})$

store a list of locally close vertices for v :
 $\mathcal{L}_v \leftarrow \{u_{q_1}, \dots, u_{q_t}\}$

Step 2.

select seed edge set:

$$E_{seeds} \leftarrow \{uv \in E(K_n) \mid u \in \mathcal{L}_v \text{ and } v \in \mathcal{L}_u\}$$

Step 3.

sort the seed edge set E_{seeds} decreasing order by edgewise

3.2.2. Growth step

The growth subroutine in this work follows the growth principals introduced in previous publications of QCM with the exception of a new parameterized method we introduce via the parameter τ . The method resolves the issue that occurs when more than one vertex is a numerically valid candidate to join a growing cluster C . Specifically, the parameter τ is intended as a sort of numerical padding for the case when similarity calculations may have slight differences between optional joins, but those differences may be a “digital artifact” of the similarity matrix calculation. That is, the optimal vertex to join may have average similarity to C greater by an insignificant margin than other possible vertices. τ allows that any vertices within an acceptable threshold of optimal may join at the same time. This new method allows for correct processing of data where symmetries exist such that some groups of data should be treated equally as opposed to being arranged in a linear order of priority. Finally, we should also mention that τ is not intended as a tuning parameter, but rather as a user determined value. We have found that $\tau = 0.008$ has performed well in all data sets tested for similarity functions where the range is in the interval $[0, 1]$.

Growth step algorithm:

Input: $E_{seeds} = \{e_{r_1}, \dots, e_{r_k}\}$ with $\{r_1, \dots, r_k\} \subset \{1, \dots, m\}$, and $S(e_{r_i}) \geq S(e_{r_{i+1}})$ for all i .
 K_n a complete graph on n vertices, and edgewise function $S : E(K_n) \rightarrow \mathbb{R}^+$,
 when AQCM is in the first iteration K_n and S are the inputs to AQCM, in
 subsequent iterations K_n and S are objects created by the previous iteration
 contraction step. For notation below we abbreviate $V(K_n)$ as just V .

Output: $\mathcal{C} = \{C_1, \dots, C_k\}$, a family of subsets $C_i \subset V(K_n)$ where each C_i is an approximation
 of some cluster (in the current graph) of relatively high edgewise density as in
[Definition 3.1](#)

Initialize: $\mathcal{C} \leftarrow \{\}$
 $\tau \leftarrow 0.008$

Step 1. If $E_{seeds} = \{\}$ **exit growth step algo** , **Else do:**

start a new cluster from the maximum weight seed:
 $C \leftarrow \{u, v\}$ where $uv = e$ is the first seed in the ordered list E_{seeds}

Step 2. try to grow cluster C :

compute α value:
 $\alpha \leftarrow 1 - \frac{1}{2(|C|+1)}$

compute $\text{den}(Q)$ where Q is the subgraph of K_n induced by C :
 see [Definition 3.1](#)

compute maximum contribution:
 see [Definition 3.2](#)
 $q \leftarrow \max(\{\text{cont}(v, Q) \mid v \in V \setminus C\})$

If $q \geq \alpha * \text{den}(Q)$ **do**:

find eligible vertices to join C :
 $J \leftarrow \{v \in V \setminus C \mid \text{cont}(v, Q) > q - \tau\}$

expand cluster C and continue:
 $C \leftarrow C \cup J$

return to Step 2.

Else do:

store cluster C :
 $\mathcal{C} \leftarrow \mathcal{C} \cup \{C\}$

eliminate unneeded seed edges and continue:
 $E_{\text{seeds}} \leftarrow E_{\text{seeds}} \setminus E(Q)$ where Q is the subgraph of K_n induced by C
return to Step 1.

The lower bound of the sequence α .

In Step 2 of the “growth step algorithm”. There is a function $\alpha \times \text{den}(Q)$ that controls the growth of an existing “cluster” Q . Some natural questions are: Is the function $\alpha \times \text{den}(Q)$ decreasing? Does it have a lower bound? If it approaches to zero in some case, then this method would fail since the growth substep will extend a cluster Q to the entire graph K_n . These questions are answered in [3]. The function $\alpha \times \text{den}(Q)$ is decreasing in general. However, the following lemma from [3] dismisses all the worries mentioned above.

Lemma ([3]). *In Step 2 of the “growth step algorithm”, the density of any new cluster Q' is at least $\frac{2}{3} \text{den}(Q_0)$ where Q_0 is the density the seed edge in the initial stage of this step.*

3.2.3. Adjustment step

First iteration adjustment algorithm:

Input: $\mathcal{C} = \{C_1, \dots, C_k\}$, a family of subsets $C_i \subset V(K_n)$ where each C_i is an approximation of some cluster (in the current graph) of relatively high edgeweight density as in [Definition 3.1](#). Observe that because it is the first iteration $C_i \subset X$, where X is the original data input.

Output: $\{C_{\ell_1}, \dots, C_{\ell_s}\} \subseteq \{C_1, \dots, C_k\}$, a clustering such that for any pair C_{ℓ_i} and C_{ℓ_j} their intersection has $|C_{\ell_i} \cap C_{\ell_j}| \leq 0.5 * \min(|C_{\ell_i}|, |C_{\ell_j}|)$

Initialize: $t \leftarrow 1$, this index tracks the current cluster.

Step 1.

sort the clusters in decreasing order by edgeweight density:
 $\mathcal{C} = \{C_{\ell_1}, \dots, C_{\ell_k}\}$ and induced subgraphs $\mathcal{Q} = \{Q_{\ell_1}, \dots, Q_{\ell_k}\}$ have
 $\text{den}(Q_{\ell_i}) \geq \text{den}(Q_{\ell_{i+1}})$ for all i .

initialize current cluster:
 set C_{ℓ_t} (t was initialized above)

Step 2. If all clusters have been checked, **exit adjustment step algo, Else do**:

find clusters overlapping current cluster too much:
 $\mathcal{H} \leftarrow \{C_{\ell_z} \mid z > t \text{ and } |C_{\ell_z} \cap C_{\ell_t}| > 0.5 * \min(|C_{\ell_z}|, |C_{\ell_t}|)\}$

Step 3. If $\mathcal{H} = \{\}$ **do**:

advance position and continue:
 $t \leftarrow t + 1$
 set current cluster as C_{ℓ_t}
goto Step 2.

Else

resolve degenerate overlap by removal:

$$\mathcal{C} \leftarrow \mathcal{C} \setminus \mathcal{H}$$

advance position and continue:

$$t \leftarrow t + 1$$

set current cluster as C_{ℓ_t}

goto Step 2.

Subsequent iterations adjustment algorithm:

Input: $\mathcal{B} = \{B_1, \dots, B_k\}$, a family of subsets $B_i \subset X$ where each B_i is an approximation of some cluster (in the original input data) of relatively high edgeweight density as in [Definition 3.1](#). Each B_i is obtained from a $C_i \in \mathcal{C}$, the output of the growth step, by associating members of $C_i \subset V(K_n)$ with the objects they represent in lower hierarchical levels.

Output: $\{B_{\ell_1}, \dots, B_{\ell_s}\} \subseteq \{B_1, \dots, B_k\}$, a clustering such that for any pair B_{ℓ_i} and B_{ℓ_j} their intersection has $|B_{\ell_i} \cap B_{\ell_j}| \leq 0.5 * \min(|B_{\ell_i}|, |B_{\ell_j}|)$

Initialize: $t \leftarrow 1$, this index tracks the current cluster.

Step 1.

sort the clusters in decreasing order by edgeweight density:

$\mathcal{B} = \{B_{\ell_1}, \dots, B_{\ell_k}\}$ and induced subgraphs $\mathcal{Q} = \{Q_{\ell_1}, \dots, Q_{\ell_k}\}$ have $\text{den}(Q_{\ell_i}) \geq \text{den}(Q_{\ell_{i+1}})$ for all i . NOTE that the subgraphs \mathcal{Q} are subgraphs of the original graph input to AQCM with vertices X and edgeweights S .

initialize current cluster:

set B_{ℓ_t} (t was initialized above)

Step 2. If all clusters have been checked, exit adjustment step algo, Else do:

find clusters overlapping current cluster too much:

$$\mathcal{H} \leftarrow \{B_{\ell_z} \mid z > t \text{ and } |B_{\ell_z} \cap B_{\ell_t}| > 0.5 * \min(|B_{\ell_z}|, |B_{\ell_t}|)\}$$

Step 3. If $\mathcal{H} = \{\}$ do:

advance position and continue:

$$t \leftarrow t + 1$$

set current cluster as B_{ℓ_t}

goto Step 2.

Else

resolve degenerate overlap by merging:

find optimal cluster for merging with current cluster:

$B^* \leftarrow B_{\ell_t} \cup B_x$, where $B_x \in \mathcal{H}$ is the cluster such that Q^* , the subgraph induced by $B_{\ell_t} \cup B_x$, has $\text{den}(Q^*)$ highest among all possible choices of $B_x \in \mathcal{H}$

update the clustering:

$$\mathcal{B} \leftarrow \mathcal{B} \setminus \{B_{\ell_t}, B_x\}$$

$$\mathcal{B} \leftarrow \mathcal{B} \cup \{B^*\}$$

re-sort the clustering and continue:

NOTE that we only need to find the correct index for placement of B^* to preserve $\text{den}(Q_{\ell_i}) \geq \text{den}(Q_{\ell_{i+1}})$ for all i

goto Step 2.

3.2.4. Contraction step

Graph contraction algorithm:

Input: $\{C_{\ell_1}, \dots, C_{\ell_s}\}$ the output of the adjustment step (equivalently $\{B_{\ell_1}, \dots, B_{\ell_s}\}$ if the iteration level of AQCM is 2 or higher, in this section we will simplify the notation by using the symbol C only). Also the graph contraction requires the original similarity data S that was input to AQCM in the first iteration.

Output: K_s a complete graph on s vertices where each vertex represents a cluster in $\{C_{\ell_1}, \dots, C_{\ell_s}\}$ and edgeweight function $S' : E(K_s) \rightarrow \mathbb{R}^+$ where $S'(C_{\ell_i}, C_{\ell_j})$ is the average similarity (according to S) between the data in C_{ℓ_i} and the data in C_{ℓ_j} .

Method:

The definition of K_s and S' are clear as described above so here we only need to define the case when $C_{\ell_i} \cap C_{\ell_j}$ is non-empty. The definition we state here is general in that for the case where $C_{\ell_i} \cap C_{\ell_j}$ is empty, the definition is equal to the statement of the output described above.

edge set definitions: We need the following definitions of edge sets of the original graph on the data X with edge set E . For C_i and C_j with $i, j \in \{\ell_1, \dots, \ell_s\}$ we have:

$$E_\beta := \{uv \in E \mid u \in C_i \text{ and } v \in C_j \text{ and } u, v \notin C_i \cap C_j\}$$

$$E_\gamma := \{uv \in E \mid u \in C_i \text{ and } u \notin C_j \text{ and } v \in C_i \cap C_j\}$$

$$E_\delta := \{uv \in E \mid u \notin C_i \text{ and } u \in C_j \text{ and } v \in C_i \cap C_j\}$$

$$E_\sigma := \{uv \in E \mid u, v \in C_i \cap C_j\}$$

edgeweight function: For C_i and C_j with $i, j \in \{\ell_1, \dots, \ell_s\}$ we define S' :

$$E_{i,j} := E_\beta \cup E_\gamma \cup E_\delta \cup E_\sigma$$

$$S'(C_i, C_j) := \frac{\sum_{e \in E_{i,j}} S(e)}{|E_{i,j}|}$$

4. Automatic cluster/community selection

The following algorithm for automatic cluster/community selection is based on the principal that a “most significant” clustering within the many possibilities encoded by a hierarchy tree should have clusters in which the internal similarity is significantly higher than that of their parent clusters [6]. Within the context of a hierarchy tree T with parent node x and child node y (representing candidate clusters C_x and C_y respectively) this change in cluster density between levels can be expressed as “density drop” $\text{den}(C_x) - \text{den}(C_y)$ (see Section 3.2 for definition of density). Furthermore, among possible clusterings exhibiting a significant density drop, we prefer smaller cluster size as we are interested in a larger number of clusters when such a clustering is an option.

The algorithm featured here assigns an edgeweight to the edge $xy \in E(T)$ that captures the properties of density drop and cluster size and then selects a clustering as an edge cut across which the average of the edgeweights is optimal among all edge cuts in the tree T . The algorithm for solving the optimal average weight cut problem is seen in Substeps 3.1–3.4 and proof of optimality is provided in [8,9]. The graph theoretic operation “edge contraction” plays an important role (see Substep 3.4 (i)) so we clarify the notation and definition here. In the directed tree T , the tree “ T contracted by e ” is notated as T/e and is constructed by: (1) delete the edge $e = xy$; (2) the vertices x and y are merged to one vertex, that is, all in edges of x and out edges of y are now in and out edges respectively of the new vertex.

Input. The generalized hierarchy tree T (on the data set X) with the root vertex v_0 , the size of every cluster C (corresponding to each node of T), the similarity density $\text{den}(C)$ of every cluster C (see Definition 3.1).

Output. $\mathcal{C} = \{C_1, \dots, C_s\}$ a clustering of the data set X .

Denote the set of leaves of T by L .

Step 1. For each directed edge $e = xy \in E(T - L)$, define the edge weight

$$w(xy) = \frac{|C_y|}{\text{den}(C_y)^2 - \text{den}(C_x)^2} \quad (1)$$

where C_x and C_y are clusters in X corresponding to the nodes x and y in T , respectively. That is, x is the parent node of the child node y .

Step 2. Find a **maximum** rooted spanning tree T_{max} of T with respect to the weight w where v_0 is the root of T_{max} .

Step 3. In this step, we are to find an edge cut Q of T_{max} separating the root v_0 and the set L of leaves such that $\frac{\sum_{e \in Q} w(e)}{|Q|}$ is **minimum** among all such edge cuts.

Substep 3.1. Calculate α_0

$$\alpha_0 = \frac{w(E^+(v_0))}{|E^+(v_0)|}.$$

Substep 3.2. For each $e \in E(T_{max} - L)$ calculate $\lambda(e_i)$ the **contractibility of the edge** e_i defined as

$$\lambda(e_i) = \frac{w(E^+(e_i)) - w(e_i)}{|E^+(e_i)| - 1}$$

and $E^+(e_i)$ is the set of all out-edges of the head of the edge e_i in the directed tree T_{max} .

Substep 3.3. Sort the edges e_i of the $E(T_{max} - L)$ so that

$$\lambda(e_1) \leq \lambda(e_2) \leq \dots \leq \lambda(e_m)$$

Substep 3.4. If $\lambda(e_1) < \alpha_0$ then

- (i) Denote the *in edge* to e_1 by e^* . Contract $T_{max} \leftarrow T_{max}/e_1$, and
- (ii) update λ value for e^* , or update α_0 if e_1 had no *in edge* (it was in $E^+(v_0)$), and
- (iii) repeat Substep 3.3.

If $\lambda(e_1) \geq \alpha_0$ then go to the final step.

FINAL STEP. In the resulting (smaller) tree T_{max} , let $Q = E^+(v_0)$. The output is

$$\{C_x \mid v_0 x \in Q, C_x \text{ is the cluster corresponding to } x \text{ in } T_{max}\}$$

We make an important note here regarding implementation of the above cluster selection algorithm. Recall that at the first iteration of AQCM it is not guaranteed that all data have joined a cluster. Indeed, at a given level in the hierarchy tree there may be some branch that represents a still unclustered data as of that level. We find that it is useful to remove such branches from the tree T_{max} before solving the optimization problem described in Step 3 above. In so doing we ensure that the resulting edge cut is related to significant drops in density of clusters between hierarchical levels without the need to define such a drop in the case of unclustered data. Furthermore this heuristic would seem relevant since it is the change in density of clusters that underpins the density drop cluster selection as a strategy.

5. Hierarchical clustering in Facebook network data

5.1. The data set

In order to test the ability of AQCM method to perform real world unsupervised learning at a scale that is large relative to the expected size of clusters we chose social network data for the test data set. The data are the Facebook friendships recorded as a “snapshot” of the state of the Facebook network at UC Berkeley in 2005. This choice of data satisfies our target of “scale relative to expected cluster size” because of the circumstances under which the network was formed. In 2005 Facebook was emerging as a prominent new social media platform at a time when other platforms such as Friendster and Myspace had already paved the way for social media to become a common aspect of young Americans’ internet use. At the time Facebook required users to sign up with an academic (university, college, or high school) email address as the platform was originally designed as an internet service to facilitate social interaction on campuses and it was within these campuses that Facebook networks first formed. These networks of student friendships represented relatively new social relationships forming around the types of social interactions fostered by campus life such as social events, fraternity/sorority life, classroom/degree related relationships. Moreover, college students tend to develop relatively small “close friend social circles” which are embedded within larger social circles related to campus activity. Hence in Facebook network data sets taken from colleges we might expect there to exist hierarchies of network community/cluster structures where on the most local level edge dense community/cluster structures are quite small relative to the network of the entire body of students on a given campus. The UC Berkeley data set represents 22,900 individuals so we might expect there to exist thousands of small clusters of close friends, thus the data set presents the type of challenge which AQCM was designed to address in automatic cluster detection.

5.2. Method of applying AQCM

5.2.1. Similarity

The UC Berkeley Facebook network is not a weighted network data set. That is, for individuals i and j who are friends the adjacency matrix has $A[i, j] = 1$ and if they are not friends $A[i, j] = 0$. Thus in order to apply AQCM it is necessary




Fig. 3. Iterative application of AQCM and edge-cut cluster/community detection.

to define similarity between two network nodes/vertices. For this definition we adopt the method presented in [25]. Mathematical details of the method are given in the appendix (Appendix A.2). The method defines for each vertex a diffusion pattern vector representing the way the vertex impacts the rest of the network in a diffusion process. Similarity between two vertices is then defined as the similarity of their diffusion vectors using “cosine of the angle between the two vectors”, a method common in data mining [7].

5.2.2. Initial results and further iterative approach

We applied AQCM to build a hierarchy tree from the similarity described above and then used edge-cut clustering to obtain 4445 clusters covering approximately 15,000 of the 22,900 vertices of the Facebook network. The hierarchy tree output by AQCM has 86,430 vertices and has 15 hierarchical levels. Clearly such a complex hierarchical clustering output presents challenges for visualization and interpretation of the clusterings it describes. These challenges motivated the development of a further iterative process within which AQCM and edge-cut clustering serve as subroutines. The iterative process is detailed in Fig. 3. It creates a simplified hierarchy tree based on the output of AQCM and edge-cut clustering run subsequently in a manner similar to AQCM itself. The method also applies some simple post-processing algorithms to clustering C (details provided in Appendix A.3) to remove multimembership and force unclustered data into clusters in a refined clustering C' . The purpose of post processing is so that diagrams may more easily describe the data set and validate the existence of hierarchical structures. For the contraction step in this iterative process we use the same contraction described in Section 3.2.4. The resulting similarity data S' is input to AQCM and the process repeats. We were able to obtain a simplified hierarchy tree suitable for diagrams in three iterations of this process. Importantly, the run time for the iterative method was not significantly longer than simply running AQCM with edge cut clustering. This is because the large share of time is spent reducing the job from 22,900 data down to roughly 5000 data (the size of the first clustering) at which point the job size is no longer large so the following iterations finish relatively quickly.

Table 1
Cluster size statistics.

Level	μ	σ	Med	Min	Max
4	602.63	657.07	337.50	46	2490
3	61.56	55.99	44	8	501
1	4.74	7.2	3	2	259




Fig. 4. (Left) Histograms of cluster sizes at each of the indicated levels. (Right) Zoom-in of the non-outlier cluster sizes at level 1.

5.3. Assessment of discovered hierarchical structures

5.3.1. Similarity and connectivity

The hierarchy tree output by the process described in Section 5.2.2 has five significant hierarchical levels and is pictured in Fig. 5-a (level numbers illustrated in red) with the similarity data described in Section 5.2.1 seen underneath as indexed by the tree. For the purpose of demonstrating the validity of hierarchical structures discovered in the UC Berkeley network we focus on levels 1, 3, and 4 as these provide a sufficient view of the scale of the clusters at higher and lower levels. These three levels feature clusterings of 4800, 372, and 38 clusters respectively. Note that here we refer to non-trivial clusters, at level 1 there are 134 un-clustered points. Some statistics on cluster size are provided for each of these three levels in Table 1 and histograms of cluster size for each of these levels are provided in Fig. 4. Level 1 has a large number of small clusters with sizes 2 and 3 accounting for 2599 of the clusters and sizes 4, 5, and 6 accounting for 1552 of the clusters. There are 315 level 1 clusters in size range 10 through 50 and only 14 with size greater than 50 with only five of those having size greater than 100.

In order to provide deeper perspective on the complex structures found in the similarity data and their relationship to the social structure of the Facebook network we follow an exploratory approach along a particular branch of the hierarchy tree. The branch represents level 4 cluster number 32 with 2490 members and is indicated with a black dot in Fig. 5-a while the internal similarity data of the cluster is highlighted by the blue box Fig. 5-a. The internal similarity is relatively high compared with the cluster's relatively low similarity to the rest of the data and visual inspection of Fig. 5-a indicates the cluster is distinguished in that way within the set of level 4 clusters. Fig. 5-e shows a depiction of the connectivity of the Facebook network with level 4 clusters illustrated by darker edges and cluster numeric labels. We computed the modularity [26] of each of the level 4 clusters and found cluster 32 has the highest modularity. Finally, accessing the metadata available, we note that 80% of cluster 32 are members of the freshman class and of the freshman class 57% are in cluster 32. Thus we can summarize that cluster 32 is suitably interesting as it has distinctness in its similarity and modularity properties and also correlates significantly with the social substructure induced by the freshman class. The existence of a distinct social subnetwork correlating with some active subset of the freshman class fits with our theory that the Facebook network should correspond with aspects of campus life. Indeed, we expect that freshmen seeking to establish a social life in their new environment would use Facebook friendships to facilitate social activity and thus a freshman oriented subnetwork should emerge.

Fig. 5-b shows a zoom in of the similarity data inside cluster 32 with the corresponding branch of the hierarchy tree. At this zoom level the lower level small clusters of high similarity begin to become visible as blocks along the diagonal

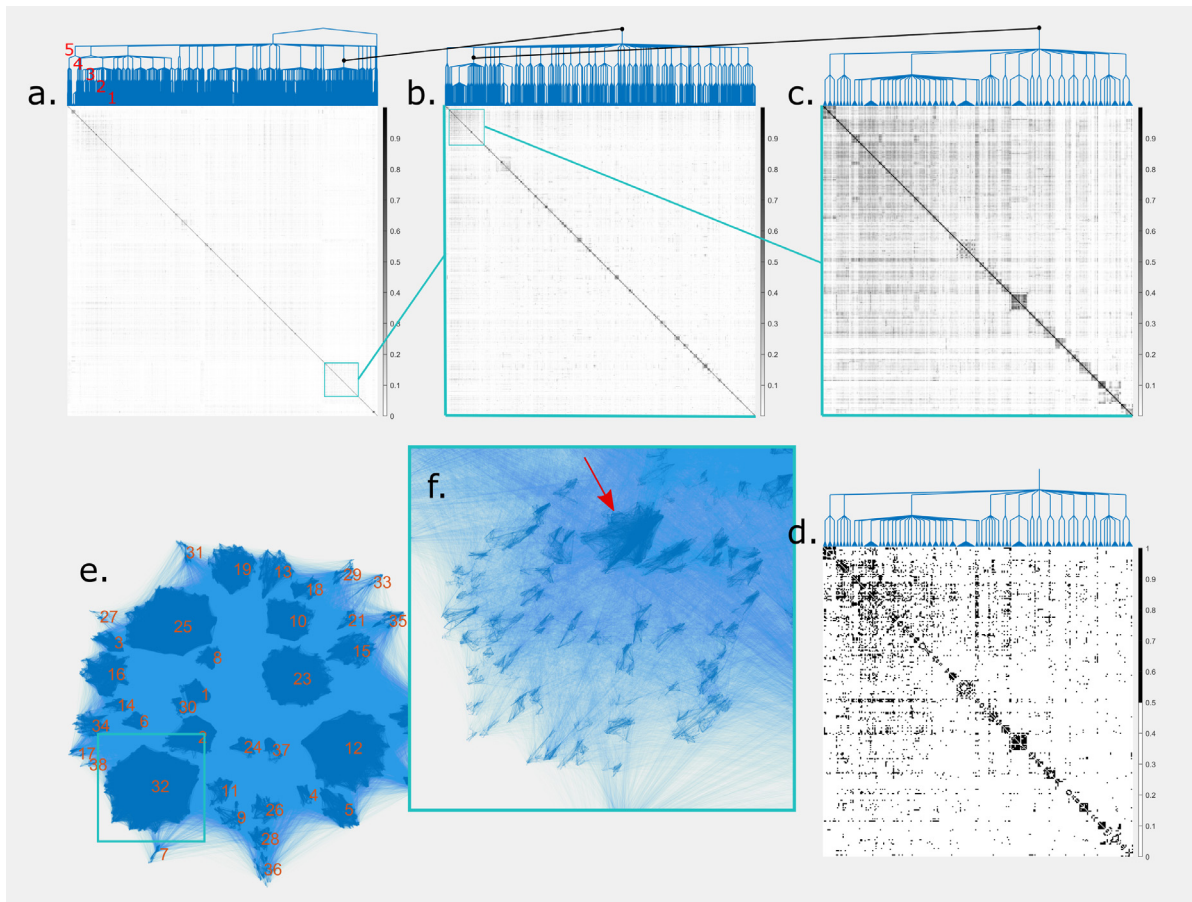


Fig. 5. Similarity and Connectivity: **a.** The hierarchy tree output by the iterative process discussed in Section 5.2.2 with five hierarchical levels labeled in red, the similarity data is shown as a matrix indexed by the tree. **b., c.** Successive zoom-ins of chosen subsets as discussed in Section 5.3.1. **d.** Same zoom-in as in c above but the data displayed is the adjacency (connectivity) data. **e.** Graphical layout of the UC Berkeley Facebook network with level 4 clusters shown darker and with cluster numeric labels. **f.** Zoom-in of the inset blue box in e, here dark edges illustrate clusters at level 3, the red arrow indicates level 3 cluster number 23 discussed in Section 5.3.1.

of the matrix. Fig. 5-f shows a zoom in of the area in the inset blue box in Fig. 5-e. In Fig. 5-f the dark edges illustrate the clusters at level 3. The red arrow indicates level 3 cluster number 23. This cluster stands out as being significantly larger. The cluster's internal similarity is highlighted by the inset blue box in Fig. 5-b. Visual inspection of Fig. 5-b reveals the relatively higher amount of similarity inside cluster 23. Interestingly cluster 23 also has relatively high similarity to many of the other clusters in Fig. 5-b,f. This property indicates that cluster 23 might be a high centrality cluster representing a hub within the Facebook network of cluster 32. In order to investigate this conjecture we compute the node degree (number of edges at a node) within cluster 32 for each member. For members of cluster 23 the average degree inside cluster 32 is 40 with a maximum of 184 while members of cluster 32 not in cluster 23 had average degree 21 with a maximum of 134. Furthermore, while cluster 23 members make up only 11% of cluster 32, more than half of "higher degree" (degree more than 50) cluster 32 members are in cluster 23. It is noteworthy though that cluster 23 is not entirely characterized as a group of high degree nodes. We observe that the average degree within cluster 32 is 24 and cluster 23 has 58% with degree greater than this, hence the degree distribution skew for cluster 23 is not very large. However, cluster 23 has 23 members with degree at least 100 while members of cluster 32 not in cluster 23 account for only 4 with degree at least 100. Hence we may characterize that cluster 23 contains almost all of the high centrality members within cluster 32 and the high centrality of a select few members of cluster 23 appears to be a key feature of the subnetwork defined by cluster 23.

Fig. 5-c,d shows the zoom-ins of the internal similarity data and adjacency data respectively of cluster 23. The adjacency data shows a complex system of smaller edge-dense submodules connected together by hub-like connectivity emanating from several clusters seen at top left. Hence the characterization of cluster 23 described above appears accurate on inspection. In the analysis below we further confirm the hub-like structures in the hierarchical subnetworks of clusters 32 and 23.

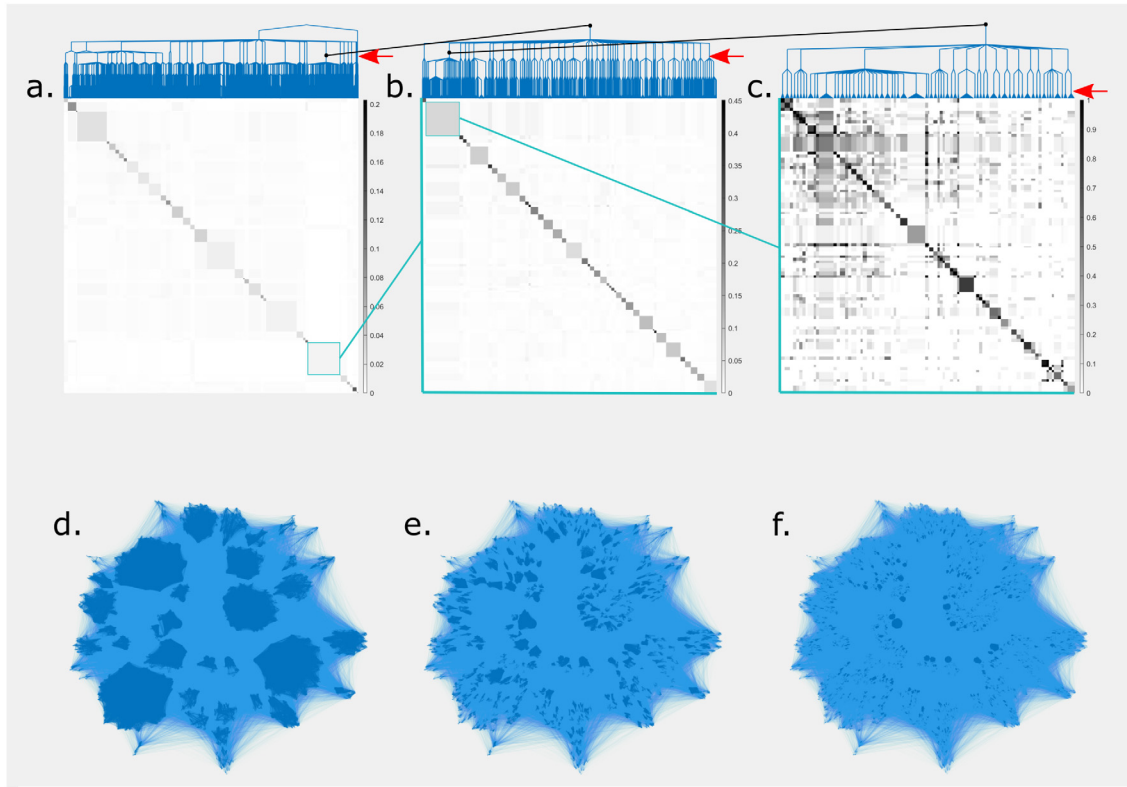


Fig. 6. *modular properties:* **a. b. c.** Shown are edge probabilities (see Section 5.3.2) within and between clusters at each of levels 4, 3, and 1 respectively. Red arrows show the hierarchical level at which edge probabilities were calculated. The zoom-ins match those in Fig. 5-a,b,c. **d. e. f.** Shown are graphical layouts of the connectivity of the UC Berkeley Facebook network with vertex placement matching Fig. 5-e. Darker edges illustrate the in-cluster edges at levels 4, 3, and 1 respectively.

5.3.2. Modular properties and hub-like structures

For the purpose of understanding the complex relationships within and between clusters we quantify the modular properties of clusters using a stochastic block model approach. For a graph G with a clustering $\mathcal{C} = \{C_1, \dots, C_s\}$ of the vertex set V the probability of an edge between clusters C_ℓ and C_k is defined as the total number of such edges divided by the possible number of such edges. If $\ell = k$ then we count the number of edges inside C_ℓ and divide by the possible number of such edges. For notation we use $P_{\ell k}$. Also note that $P_{\ell k}$ can equivalently be referred to as edge density. Community structures in networks are understood to be relatively edge-dense subsets of vertices separated from the rest of the network by a relatively sparse edge cut. Under this model we would say a cluster (community) C_ℓ should have $P_{\ell\ell} > P_{\ell x}$ where C_x is defined as the vertex set $V \setminus C_\ell$. In this work we are interested in all the values $P_{\ell k}$ over the clustering \mathcal{C} . For visual analysis we define a matrix $M_{\mathcal{C}}$ based on these values. For vertices $v_i \in C_\ell$ and $v_j \in C_k$ we assign $M_{\mathcal{C}}[i, j] = P_{\ell k}$. We computed the matrix $M_{\mathcal{C}}$ for each of the three clustering levels 1, 3, and 4. $M_{\mathcal{C}}$ for level 4 is seen in Fig. 6-a as indexed by the hierarchy tree above it. The red arrow indicates level 4. Figs. 6-b,c show $M_{\mathcal{C}}$ for levels 3 and 1 respectively and zoomed in the same way as in Figs. 5-b,c.

In order to assess the distinctness of the clusters in terms of their modular properties as described above we may consider, for a cluster C_ℓ , the ratios of the form $(P_{\ell\ell} - P_{\ell k})/P_{\ell\ell}$ over all choices of $k \neq \ell$. For a cluster C_ℓ we will notate this ratio as $\Delta_\ell(k)$. Furthermore it is useful to consider, for a cluster C_ℓ , the minimum over k of $\Delta_\ell(k)$. For this minimum we use the notation $\delta_\ell(k)$.

Using the framework described above we first observe that for the level 4 clustering $\Delta_\ell(k)$ is positive for each C_ℓ with an average of 94%. The average of $\delta_\ell(k)$ over all clusters is 83.8% with standard deviation 11.8%. Also, the minimum value of $\delta_\ell(k)$ observed is 51.9% indicating that even in the case where two clusters are not very distinctly separated there is still a significantly lower probability of between cluster edges. Hence the level 4 clustering exhibits significant community structure in that the difference in internal and between cluster edge probability is statistically significant relative to cluster edge density. In particular we have for cluster 32 the average over k of $\Delta_{32}(k)$ is 88.9%. For 86.5% of the k values $\Delta_{32}(k)$ is at least 75% validating that cluster 32 is a strongly separated community, a property that agrees with its high modularity as mentioned in Section 5.3.1. We may also observe that visual inspection of Fig. 6-a reveals that the edge probabilities between cluster 32 and most other clusters in level 4 is significantly lower than between cluster probabilities among other clusters, further demonstrating that cluster 32 stands out as a particularly well separated subnetwork.

Repeating the analytic inquiry above for the clustering at level 3 we find the following. $\Delta_\ell(k)$ is positive for each C_ℓ with an average of 98.9%. The average of $\delta_\ell(k)$ over all clusters is 85.4% with standard deviation 7.9%. Also, the minimum value of $\delta_\ell(k)$ observed is 38% indicating that in some cases cluster separation is not as distinct and may be on account of a cluster group forming a higher level cluster. Hence the level 3 clustering exhibits community structure distinctness similar to that at level 4. Indeed, if we explore the internal connectivity of level 4 cluster 32 as an example, the level 3 cluster to cluster edge probabilities seen in Fig. 6-b clearly show the relatively edge dense level 3 clusters interconnected by relatively sparse connectivity.

Another feature that stands out in visual inspection of Fig. 6-b is the apparently higher edge probability of edges to cluster 23 (blue inset box). While the edge probability from any subcluster of 32 to subcluster 23 averages 0.9%, the expected value of such averages for other subclusters seen in Fig. 6-b is just 0.47% and for 90% of subclusters of 32 that average is less than 0.7%. Cluster 23 of course has the highest edge probability average hence, in the sense of the statistics mentioned here, we may confirm our conjecture made in Section 5.3.1 that subcluster 23 is a hub within cluster 32.

Repeating the analytic inquiry above for the level 1 clustering is more challenging on account that 1039 of the 4800 level 1 clusters have the property $P_{\ell\ell} < P_{\ell k}$ for at least one value of k . Clusters such as this can arise when common neighborhood properties are more relevant for a set of vertices than direct connections. Consider for example ten Facebook members who are all friends with each other. Suppose there are three additional members who are not friends with each other but who are each friends with all of the ten. In this case the three would form a natural cluster of similarity but with no internal edge density in their group. In order to describe the edge dense properties of level 1 clusters we focus the statistics on those with $P_{\ell\ell} > P_{\ell k}$. There are 3087 such clusters. Among those the average of $\Delta_\ell(k)$ is 99.7% and the average of $\delta_\ell(k)$ over those clusters is 40.6%. Hence 64% of level 1 clusters are simply described by the edge-dense community model.

Inspecting Fig. 6-c it is clear that one of the most noticeable features is the apparently hub structure formed by the prominent cluster top left. This cluster is number 1123 of the level 1 clustering. The average of $P_{1123,k}$ over clusters k featured in Fig. 6-c is 17.34% which puts cluster in the top 8.9% of subclusters of cluster 23 in terms of connectivity to other clusters. At 14 members, cluster 1123 is significantly larger than other cluster 23 subclusters with hub-like connectivity, hence we may think of it as a prominent feature of the subnetwork of level 3 cluster 23 as conjectured in Section 5.3.1.

CRedit authorship contribution statement

Scott Payne: Software development, MATLAB programming, Synthetic data design, Data testing and analysis, Figure composition, Writing. **Edgar Fuller:** Organizing, Figure editing, Mathematics theory, Text editing. **George Spirou:** Organizing, Funding acquisition, Biology theory. **Cun-Quan Zhang:** Algorithm design, Supervision, Mathematics theory, Writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

A.1. Additional graph theory terminology

For a vertex v and an edge e (in an undirected graph), we say e is *incident* to v if $e = uv$ for some vertex u . The *degree* of a vertex v is the number of edges incident to v and is notated $d(v)$. If G is a weighted graph then degree is generalized as $d(v) = \sum_{u \in V} w(uv)$, where if $uv \notin E$ then $w(uv) = 0$. If v is a vertex of a directed graph we separately count the *out-degree* (the number of edges with v as a tail) and the *in-degree* (the number of edges with v as a head), and they are notated $d^+(v)$ and $d^-(v)$ respectively. Generalization of weighted in and out degree follow the same intuition as for the undirected case.

A *path* between two vertices $u, v \in V$ is a sequence $u = u_1 u_2 \dots u_k = v$ such that for each consecutive pair $i, i + 1$ we have, $u_i u_{i+1} \in E$ and each vertex appears only once in the sequence. For two specific vertices u and v a path between them is often called a u, v -path. A *closed path* is a path where the first and last vertex are the same and is often called a *cycle*. We say a graph is *connected* when for any pair of vertices $u, v \in V$ there is a u, v -path. A set of edges $H \subset E$ is called an *edge cut* when the removal of the edges H causes the graph to become disconnected. In many graph theoretic works it is important to restrict study to edge cuts that split a graph into a set of *non-trivial* components $\{Q_1, \dots, Q_k\}$. A non-trivial component Q of a graph G is a *subgraph* of G with $V(Q) \subsetneq V(G)$ and $E(Q) \subsetneq E(G)$ and $|V(Q)| > 1$ and Q is connected and importantly for $u, v \in V(Q)$ with $uv \in E(G)$ then $uv \in E(Q)$. Edge cuts that split a graph G into non-trivial components are called non-trivial edge cuts, but in context it is often understood, in a given work, that the term *edge cut* refers to those that are non-trivial.

A minimal connected graph on a vertex set V with $|V| = n$ is any graph such that no edge can be removed without disconnecting the graph. It is easy to see that any such graph has no cycle and easy to show that any such graph has $n - 1$

edges. A minimal connected graph is called a *tree* and often such a graph is notated T instead of G . An obvious property of any tree T is that for any pair of vertices u and v there is exactly one u, v -path.

In the study of directed graphs, an important type of tree is the *rooted tree*. A rooted tree is a directed tree with exactly one vertex v with $d^-(v) = 0$ (this vertex is called the *root*) and set of vertices $\{u_1, \dots, u_k\}$ with $d^+(u_i) = 0$ for each (these vertices are called *leaves*). Stated another way, for a rooted tree T , there is one vertex (root) that is the *head* of no edge and some set of vertices (leaves) that are each the tail of no edge. It is easy to show that for a root vertex v and any non-root vertex u there is a v, u -dipath (directed path), hence a rooted tree can be thought of as having a “layered” structure in that nodes can be organized into groups by their distance from the root along root-to-leaf paths.

A.2. The similarity matrix, S

We give here a brief summary of the construction of the matrix S , that we used in our analysis of the FaceBook network data. Detailed study of this construction of similarity matrix can be seen in [25].

Input. Given an adjacency matrix A of an input graph G where the (i, j) -entry $A[i, j]$ is the weight of the directed edge $v_i \rightarrow v_j$. Let $n = |V(G)|$. (If G is undirected, A is symmetric; If G is unweighted, $A[i, j] = 0$ or 1.)

Output. A diffusion similarity matrix S where the (i, j) -entry $S[i, j]$ is the similarity between the vertices v_i and v_j .

NOTE: Performing **Step 1.** for any directed graph input will save computation time needed to otherwise confirm the strong-connected property for the graph.

Step 1. If G is strongly connected, then go to Step 2. Otherwise, let v_{n+1} be a new vertex, and

$$V(G) \leftarrow V(G) \cup \{v_{n+1}\},$$

$$E(G) \leftarrow E(G) \cup \{v_i \rightarrow v_{n+1} : i = 1, \dots, n\} \cup \{v_{n+1} \rightarrow v_i : i = 1, \dots, n\}$$

and for every $i = 1, \dots, n$,

$$A[i, n+1] = A[n+1, i] = 10\% \min\{A[i, j] : A[i, j] > 0, \forall i, j = 1, \dots, n\}$$

then go to Step 2.

Step 2. Construct W where the (i, j) -entry

$$W[i, j] := \frac{A[i, j]}{\sum_{\forall \mu} A[i, \mu]} \quad (2)$$

Step 3. Construct

$$K := \sum_{k=1}^{\infty} \frac{(\approx 1.63)^k}{k!} W^k \quad (3)$$

Step 4. If $|V(G)| = n+1$ (that is, Step 1 was taken), then deleting the $(n+1)$ st-row and $(n+1)$ st-column from K . And go to Step 5.

If $|V(G)| = n$, then go to Step 5.

Step 5. Let $\vec{K}_{row}(i)$ be the i th-row of K and $\vec{K}_{col}(i)$ be the i th-column of K . Construct the output S where the (i, j) -entry

$$S[i, j] := \frac{\cos(\vec{K}_{col}(i), \vec{K}_{col}(j)) + \cos(\vec{K}_{row}(i), \vec{K}_{row}(j))}{2} \quad (4)$$

where $\cos(\vec{\alpha}, \vec{\beta})$ is the cosine of the angle between vectors $\vec{\alpha}$ and $\vec{\beta}$. **END.**

A.3. Post processing

Let X be a data set and $\mathcal{C} = \{C_1, \dots, C_s\}$ be a family of subsets of X , possibly with $|C_i| = 1$ for some C_i and possibly with some pair of subsets having non-empty intersection. Let $S(x, y)$ be the similarity value of data x and y and let the similarity density of C_i , notated $\text{den}(C_i)$, and the contribution of x to C_i , notated $\text{cont}(x, C_i)$, be as defined in Section 3.2. If $|C_i| = 1$ then we shall define that $\text{den}(C_i) = 1$. The following post processing algorithms augment the clustering \mathcal{C} so as to accommodate certain properties of clustering output that may be desirable for some applications.

A.3.1. Expanding the cluster set by including some un-clustered vertices

The AQCM algorithm builds a hierarchical description of a data set based on preserving local similarity density. It is possible for a given data set that the output clustering of AQCM and cluster detection will leave some data points unclustered as “outliers” of a significantly more locally dense selected cluster. For various data analytic purposes it may be desirable to place these unclustered data into clusters. These unclustered data also may naturally form small clusters themselves if their local cluster was detected at a position in the hierarchy tree higher than a cluster detection cut for some nearby denser cluster. For these reasons we have developed the following post processing algorithm which allows

unclustered data to naturally “self-assign” to another unclustered data or a nearby cluster based on a factor described below.

The following definitions quantify the similarity of data x to C_i as relative to other clusters and also as relative to the similarity density of C_i . In this way the *clustering factor* $\phi_c(x)$ is a value that represents a best fit of data x to a cluster in that data x should be most similar to its ideal cluster without reducing the cluster’s similarity density too much.

Definition A.1. For a given $x \in X$ and $C_i \in \mathcal{C}$ with $C_i \neq \{x\}$,

(1) The *individual preference* of x to C_i .

$$\phi_p(x, C_i) = \frac{\text{cont}(x, C_i)}{\max\{\text{cont}(x, C_j) \mid C_j \in \mathcal{C} \text{ with } C_j \neq \{x\}\}}$$

(2) The *community acceptance* of x by C_i .

$$\phi_a(x, C_i) = \frac{\text{cont}(x, C_i)}{\text{den}(C_i)}$$

(3) The *mutual preference* between x and C_i .

$$\phi_m(x, C_i) = \phi_p(x, C_i) \times \phi_a(x, C_i)$$

(4) The *clustering factor* of x ,

$$\phi_c(x) = \max\{\phi_m(x, C_i) \mid C_i \in \mathcal{C}\}.$$

A.3.2. Expansion algorithm

Input: $\mathcal{C} = \{C_1, \dots, C_s\}$ a clustering of a data set X with $|C_i| \geq 2 \ \forall i$, and a threshold parameter ρ .

Output: $\mathcal{C}' = \{C_1, \dots, C_t\}$, with $t \geq s$ and \mathcal{C}' covers more of the data X than \mathcal{C} .

Step 1: List the unclustered data: $X' \leftarrow \{x \in X \mid x \notin C_i \text{ for any } C_i \in \mathcal{C}\}$.

Step 2: Update \mathcal{C} as follows: for each $x \in X'$ create a cluster $C = \{x\}$ and store the cluster $\mathcal{C} \leftarrow \mathcal{C} \cup \{C\}$.

Step 3: For each $x \in X'$ and every $C_i \in \mathcal{C}$, calculate $\phi_c(x)$ and $\phi_m(x, C_i)$.

Step 4: For each $x \in X'$, if there is a unique cluster C_i for which $\phi_c(x) = \phi_m(x, C_i)$, and $\phi_c(x) \geq \rho$, then add x to C_i .

Step 5: From \mathcal{C} , remove C_i if $|C_i| \leq 1$, and delete any duplicate cluster. $\mathcal{C}' \leftarrow \mathcal{C}$ and output \mathcal{C}' .

END

A.3.3. Elimination of multi-membership property

The growth subroutine of AQCM allows clusters to include optimally similar points regardless of whether those points have already joined another cluster to which they are also optimally similar. Thus points representing the boundary of clusters may be “multi-members”. For some analytic purposes however it may be desirable to enforce a strict partitioning of a data set into clusters. For this reason we provide the following simple method to require multi-member data points to “choose” one cluster in which to remain a member.

We may notate the multi-member set precisely. Let $X'' \subset X$ such that for $x \in X'' \ \exists \mathcal{H}_x \subset \mathcal{C}$ where $x \in C_i \ \forall C_i \in \mathcal{H}_x$ and $|\mathcal{H}_x| \geq 2$.

Definition A.2. For a given $x \in X''$ and $C_i \in \mathcal{C}$

(1) The *contribution* of x to the *core* of C_i .

$$\varphi(x, C_i) = \text{cont}(x, C_i \setminus X'')$$

(2) The *core factor* of x .

$$\varphi_c(x) = \max\{\varphi(x, C_i) \mid C_i \in \mathcal{C}\}$$

A.3.4. Multi-membership elimination algorithm

Input: $\mathcal{C} = \{C_1, \dots, C_s\}$ a clustering of a data set X and some pairs C_i, C_j have $C_i \cap C_j \neq \emptyset$

Output: $\mathcal{C}' = \{C_1, \dots, C_t\}$, with the number of multi-member data reduced.

Step 1: List the multi-member data: $X'' \leftarrow$ multi-member data points as described above

Step 2: For each $x \in X''$ and every $C_i \in \mathcal{C}$, calculate $\varphi_c(x)$ and $\varphi(x, C_i)$.

Step 3: For $x \in X''$, if there is a unique cluster C_i for which $\varphi_c(x) = \varphi(x, C_i)$ then remove x from all clusters except C_i .

Step 4: From \mathcal{C} , remove C_i if $|C_i| \leq 1$, and delete any duplicate cluster. $\mathcal{C}' \leftarrow \mathcal{C}$ and output \mathcal{C}' .

END

References

- [1] Anil K. Jain, Data clustering: 50 years beyond K-means, *Pattern Recognit. Lett.* 31 (8) (2010) 651–666, Award winning papers from the 19th International Conference on Pattern Recognition, ICPR.
- [2] R. Xu, D.C. Wunsch, Clustering algorithms in biomedical research: A review, *IEEE Rev. Biomed. Eng.* 3 (2010) 120–154.
- [3] Yongbin Ou, Cun-quan Zhang, A new multimembership clustering method, *J. Ind. Manag. Optim.* 3 (2007).
- [4] Yongbin Ou, Lan Guo, Cun-Quan Zhang, A new clustering method and its application to proteomic profiling for colon cancer, in: *IASTED International Conference on Computational and Systems Biology : November 13–14, 2006, Dallas, Texas, USA, CASB 2006, vol. 2006, 2006*, pp. 68–72.
- [5] Peixin Zhao, Cun-Quan Zhang, A new clustering method and its application in social networks, *Pattern Recognit. Lett.* 32 (15) (2011) 2109–2118.
- [6] Xingqin Qi, Wenliang Tang, Yezhou Wu, Guodong Guo, Eddie Fuller, Cun-Quan Zhang, Optimal local community detection in social networks based on density drop of subgraphs, *Pattern Recognit. Lett.* 36 (2014) 46–53.
- [7] Ali Seyed Shirshorshidi, Saeed Aghabozorgi, Teh Ying Wah, A comparison study on similarity and dissimilarity measures in clustering continuous data, *PLoS One* 10 (12) (2015) 1–20.
- [8] Scott Payne, Edgar Fuller, Cun-Quan Zhang, Edge-cuts of optimal average weights, *Asia-Pac. J. Oper. Res.* 36 (02) (2019) 1940006.
- [9] Scott Payne, Edgar Fuller, Cun-Quan Zhang, Edge-cuts optimized for average weight: a new alternative to ford and fulkerson, 2020, *ArXiv arXiv:2002.00263*.
- [10] Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, Etienne Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech. Theory Exp.* 2008 (10) (2008) P10008.
- [11] Lucas G.S. Jeub, Olaf Sporns, Santo Fortunato, Multiresolution consensus clustering in networks, *Sci. Rep.* 8 (1) (2018) 3259.
- [12] S. Payne, *Aqcm algorithm MATLAB*, 2021, https://github.com/scottpayne282/AQCM_algorithm_MATLAB.
- [13] Brendan J. Frey, Delbert Dueck, Clustering by passing messages between data points, *Science* 315 (5814) (2007) 16.
- [14] J. Pearl, Reverend Bayes on inference engines: A distributed hierarchical approach, in: *Proceedings of the Second National Conference on Artificial Intelligence*, AAAI Press, Menlo Park, California, 1982, pp. 133–136.
- [15] A. Rodriguez, A. Laio, Clustering by fast search and find of density peaks, *Science* 344 (6191) (2014) 1492–1496.
- [16] D. Cheng, Q. Zhu, J. Huang, Q. Wu, L. Yang, Clustering with local density peaks-based minimum spanning tree, *IEEE Trans. Knowl. Data Eng.* 33 (2) (2021) 374–387.
- [17] M.A.T. Figueiredo, A.K. Jain, Unsupervised learning of finite mixture models, *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (3) (2002) 381–396.
- [18] M. Girvan, M.E.J. Newman, Community structure in social and biological networks, *Proc. Natl. Acad. Sci.* 99 (12) (2002) 7821–7826.
- [19] V.A. Traag, L. Waltman, N.J. van Eck, From Louvain to Leiden: guaranteeing well-connected communities, *Sci. Rep.* 9 (1) (2019) 5233.
- [20] Anil K. Jain, Martin H.C. Law, Data clustering: A user's dilemma, in: Sankar K. Pal, Sanghamitra Bandyopadhyay, Sambhunath Biswas (Eds.), *Pattern Recognition and Machine Intelligence*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 1–10.
- [21] Leto Peel, Daniel B. Larremore, Aaron Clauset, The ground truth about metadata and community detection in networks, *Sci. Adv.* 3 (5) (2017).
- [22] Adrian Bondy, M. Ram Murty, *Graph Theory*, Springer, 2008.
- [23] D. West, *Introduction To Graph Theory*, Prentice Hall, 2001.
- [24] R. Diestel, *Graph Theory*, Springer, 2017.
- [25] Scott Payne, Edgar Fuller, George Spirou, Cun-Quan Zhang, Diffusion profile embedding as a basis for graph vertex similarity, *Netw. Sci.* 9 (3) (2021).
- [26] Mark E.J. Newman, Modularity and community structure in networks, *Proc. Natl. Acad. Sci.* 103 (23) (2006) 8577–8582.