

Clustering and Visualization of Fuzzy Communities In Social Networks

Timothy C. Havens

Department of Electrical and Computer Engineering
Department of Computer Science
Michigan Technological University,
Houghton, MI 49931 USA
thavens@mtu.edu

James C. Bezdek, Christopher Leckie, Jeffery Chan,
Wei Liu, James Bailey, Kotagiri Ramamohanarao,
Marimuthu Palaniswami
Department of Computer and Information Systems
Department of Electrical and Electronic Engineering
University of Melbourne
Victoria 3010, Australia

Abstract— We discuss a new formulation of a fuzzy validity index that generalizes the Newman-Girvan (NG) modularity function. The NG function serves as a cluster validity functional in community detection studies. The input data is an undirected graph $G = (V, E)$ that represents a social network. Clusters in V correspond to socially similar substructures in the network. We compare our fuzzy modularity to an existing modularity function using the well-studied Karate Club data set.

Keywords—fuzzy communities; community detection; modularity; fuzzy modularity; specVAT; clique discovery

I. INTRODUCTION

Suppose $O = \{o_1, \dots, o_n\}$ denotes a set of n objects, usually, but not restricted to, humans (karate students, monks, southern women, etc.). Let $R = [r_{ij}]$ be a matrix of relational values on $O \times O$, r_{ij} being the relation between o_i and o_j . A common form of R arises as dissimilarity data, say $D = [d_{ij}]$, where d_{ij} is the pair wise dissimilarity between object vectors \mathbf{x}_i and \mathbf{x}_j in \mathfrak{R}^p , $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$. In this case D is a *symmetric* matrix of distances. But for other types of (dis)similarity data, $d_{ij} = d(o_i, o_j)$ may not be symmetric, $d_{ij} \neq d_{ji}$. For example, Sampson's monastery data [1] is of this type. Breiger et al. [2] give the relationship from Bonhaven to Ambrose the value 2 in Sampson's data, but the value from Ambrose to Bonhaven in the opposite direction is 1. According to Wasserman and Faust [3] this is the most common form of social network data.

The Wasserman and Faust text is arguably the "bible" for social network analysis (18th printing, 2009), and yet, it does not mention fuzzy models of social networks! This is probably due to the well known disconnect between various communities of scholars working in related but uncommunicative fields. Selected readings in the literature from various groups indicate that this is probably quite accidental, most likely due to a lack of time to explore what may be essentially similar approaches advanced by disparate groups of researchers.

But many recent papers do exhibit fuzzy or possibilistic clusters in social networks. We will begin with a short review

of the evolution of fuzzy models in social networks that culminates with current work about overlapping (fuzzy) communities in social networks. Then we will develop a new measure of fuzzy modularity for community detection, and compare it to an existing one using Zachary's Karate Club [4] data set.

II. FUZZY MODELS FOR SOCIAL NETWORKS

Social network analysis usually begins with a crisp (meaning not fuzzy, probabilistic or possibilistic) graph-theoretic representation of the social network, say $G = (V, E, W)$, where V is the vertex set, E is the edge set, and W is the set of edge weights. Different social situations are realized by graphs with various properties: directed or not, weighted or not, connected or not, complete or not, and so on. In this note, G is undirected and weighted. Clusters (cliques, subtrees, etc.) in G (subsets of vertices in V) represent groups of individuals that are somehow related to each other more closely than to the individuals in the other clusters.

Any weighted graph can be thought of as a (possibly unnormalized) fuzzy graph, or a fuzzy similarity relation on pairs of nodes, first discussed by Zadeh in [5]. The earliest work on the use of fuzzy relations for social network analysis was Blin [6], who introduced the idea of using fuzzy relations in group decision theory. Bezdek et al. [7-9] collected data from small groups of students in communications classes, and developed models based on reciprocal fuzzy relations that quantified notions such as distance to consensus.

An idea that is gaining traction in social network analysis is the notion of overlapping communities in social networks [10]. Communities are defined as groups of densely interconnected nodes that are only loosely connected to the rest of the network in [11]. There is no clustering algorithm in [11]. Instead, overlapping clusters are seen visually as off-diagonal content in co-appearance images of the connection data. The model in [12] finds fuzzy communities by *multicut* spectral clustering. Clustering is done by both hard/fuzzy c-means (HCM/FCM, [13]) and validation is done with an index called fuzzy modularity by the authors.

III. PARTITIONS AND MODULARITY

Clustering in unlabeled data is the assignment of *labels* to the objects in O . Let (c) be an integer, $1 \leq c \leq n$. A *c-partition* of X is a set of (cn) values $\{u_{ik}\}$ arrayed as a $c \times n$ matrix $U = [u_{ik}]$. Element u_{ik} is the *membership* of o_k in cluster i . There are three sets of partition matrices [13]:

$$M_{pcn} = \left\{ U \in \mathfrak{R}^{c \times n} : 0 \leq u_{ik} \leq 1 \forall i, k; \sum_{i=1}^c u_{ik} \leq c \forall i; 0 < \sum_{k=1}^n u_{ik} < n \forall k \right\}; \quad (1)$$

$$M_{fcn} = \left\{ U \in M_{pcn} : \sum_{i=1}^c u_{ik} = 1 \forall k \right\}; \quad (2)$$

$$M_{hcn} = \left\{ U \in M_{fcn} : u_{ik} \in \{0, 1\} \forall i, k \right\}. \quad (3)$$

Equations (1-3) define, respectively, the sets of non-degenerate (no row is all zeroes) possibilistic, constrained fuzzy or probabilistic, and crisp c -partitions of X . Note that $M_{hcn} \subset M_{fcn} \subset M_{pcn}$.

Each run of a clustering algorithm on any data set produces one or more U 's in some M_{pcn} . For example, fixing all control and model parameters except c , applying any c -means algorithm to X produces one $U(c)$ in M_{pcn} for each $c = 2, 3, \dots, n-1$. Other runs with variations of the c -means parameters, or other clustering algorithms, produce other U 's for consideration. We collect all the *candidate partitions* into a set named CP , and ask: which $U \in CP$ is the most satisfactory explanation of substructure in O ? This is the cluster validity (more simply, "validation") problem [13].

Social network data are represented by a graph $G = (V, E, W)$, where V is a set of n vertices. E is a set of m edges, and W is a set of edge weights. Clustering in the graph G means finding partitions $U \in M_{pcn}$ of V . Because the data are not object vectors or dissimilarity data, as is usually the case in cluster analysis, finding candidate partitions often requires special methods. And the derivative problem of validating the found clusters (cluster validity) for this special type of data structure is also treated somewhat differently than the validation schemes often employed in pattern recognition.

Validity indices for partitions of G are usually called *quality functions* in the community detection literature. According to Fortunato [10]:

A quality function is a function that assigns a number to each partition of a graph. In this way one can rank partitions based on their score given by the quality function. Partitions with high scores are "good," so the one with the largest score is by definition the best. ... The most popular quality function is the modularity of Newman and Girvan [14].

The basic rationale for modularity is that a random graph doesn't have cluster structure, so the existence of clusters is revealed by comparing the actual density of edges in a subgraph to the expected density under some null hypothesis. The expected edge density depends on the chosen null model.

The most popular form of modularity assumes that W is organized as an $(n \times n)$ positive, symmetric edge weight matrix of G . Let V be partitioned into c crisp subsets of vertices (indices), say $\{V_1, \dots, V_c\}$, let $U \in M_{hcn}$ be the crisp c -partition of G . The modularity of U for $G = (V, E, W)$ is [14]:

$$Q_h = \sum_{k=1}^c \left[\frac{S(V_k, V_k)}{S(V, V)} - \left(\frac{S(V_k, V)}{S(V, V)} \right)^2 \right], \quad (4)$$

where $S(V_a, V_b) = \sum_{i \in V_a, j \in V_b} w_{ij}$ and the subscript h attached to Q

indicates that U is crisp (hard). A second equivalent form of (4) given in Fortunato [10] follows by letting

$m_i = \sum_{k=1}^n w_{ik} = S(\{i\}, V)$ and $\|W\| = \sum_{j=1}^n \sum_{i=1}^n w_{ij}$. Then (4) becomes

$$Q_h = \frac{1}{\|W\|} \sum_{k=1}^c \left(\sum_{i,j \in V_k} \left[w_{ij} - \frac{m_i m_j}{\|W\|} \right] \right). \quad (5)$$

Although the partition U takes part in the calculation of (4) or (5), its role is somewhat obscured by these forms of the modularity index. It is not hard to show that (see [15] for a proof) if the vector $\mathbf{m} = (m_1, \dots, m_n)^T = W \mathbf{1}_n$ and $B = [W - (\mathbf{m}^T \mathbf{m} / \|W\|)]$, we can also write Q_h in the more transparent form

$$Q_h(U) = \text{tr}(UBU^T) / \|W\|, U \in M_{hcn}. \quad (6)$$

Equation (6) explicitly reveals the role played by the partition U of V in the computation of modularity Q_h . The very important point about this version of modularity is that this formula is well-defined for *any* partition of V , not just crisp ones. We define the *generalized modularity* of U wrt $G = (V, E, W)$ as

$$Q_g(U) = \text{tr}(UBU^T) / \|W\|, U \in M_{pcn}. \quad (7)$$

Q_g is a proper generalization of the Newman-Girvan modularity because (7) reduces to (4) or (5) when U is a crisp c -partition of V , i.e., $Q_g|_{U \in M_{hcn}} = Q_h$. Consequently, we are entitled to call (7) the *fuzzy modularity* of U when $U \in M_{fcn}$ is a fuzzy c -partition of the vertices in V .

So far, we have not described any method for *finding* a fuzzy c -partition of V , but once we have a set CPs, we have a means for assessing the quality of each candidate in it, namely Q_g . Brandes et al. [16] review "an array of heuristic algorithms that have been proposed to optimize modularity based on greedy agglomeration, spectral division, simulated annealing and extremal optimization, to name but a few prominent examples." None of the references given in [16] uses the explicit formulation for Q_h shown in (7). We conjecture here, but leave to another study [15], the possibility that imbedding Q_h in the more general setting afforded by Q_g will lead to a new, possibly better way, to maximize this popular index.

Several other formulas that are also called fuzzy modularity appear in the literature [12, 17]. We are interested here in the formulation due to Zhang et al. [12]. Their fuzzy version of (4) begins by partitioning V with spectral clustering applied to G using FCM once the eigenvector representation of G is selected. After a fuzzy c -partition $U \in M_{fcn}$ is found this way, they convert it to a *possibilistic* c -partition $U \in M_{pcn}$ of V as follows: they choose a threshold λ , (presumably $0 < \lambda < 1$), and extract from the k -th column of $U \in M_{fcn}$ the index set $V_k = \{i \mid u_{ik} > \lambda; 1 \leq i \leq c\}$. For each vertex i in V_k , the value u_{ik} in the fuzzy c -partition is replaced by a 1. After a pass over all n columns of U is completed, the remaining (non-1) memberships are set to 0. Doing this for $k = 1$ to n results in the conversion $U \in M_{fcn} \rightarrow U_\lambda \in M_{pcn}$. Figure 1 is an example of the conversion procedure that illustrates the conversion for $\lambda = 0.10$ and $\lambda = 0.20$.

$$U = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 & p_5 \\ 0.02 & 0.12 & 0.40 & 0.44 & 0.91 \\ 0.08 & 0.18 & 0.22 & 0.56 & 0.05 \\ 0.90 & 0.70 & 0.38 & 0.00 & 0.04 \end{bmatrix}$$

↓

$$U_{0.1} = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 \end{bmatrix}; U_{0.2} = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 \end{bmatrix}$$

Figure 1. Illustration of Zhang et al.'s Conversion

Imagine that FCM is applied to the spectral data from a network of 5 people $\{p_k\}$ at $c = 3$, and terminates at the fuzzy 3-partition U shown in Figure 1. Then Zhang et al.'s conversion yields the possibilistic 3-partitions $U_{0.1}$ and $U_{0.2}$ shown below U in the figure. Zhang et al. interpret "fuzzy communities" in $U_{0.1}$ as follows. Persons '2' and '3' belong to all three groups; '4' belongs to groups 1 and 2, while '1' is only in group 3 and '5' is only in group 1.

Now consider the second partition shown in Figure 1. When λ is increased from 0.1 to 0.2, joint membership in fuzzy communities is more stringent. Now only person '3' belongs

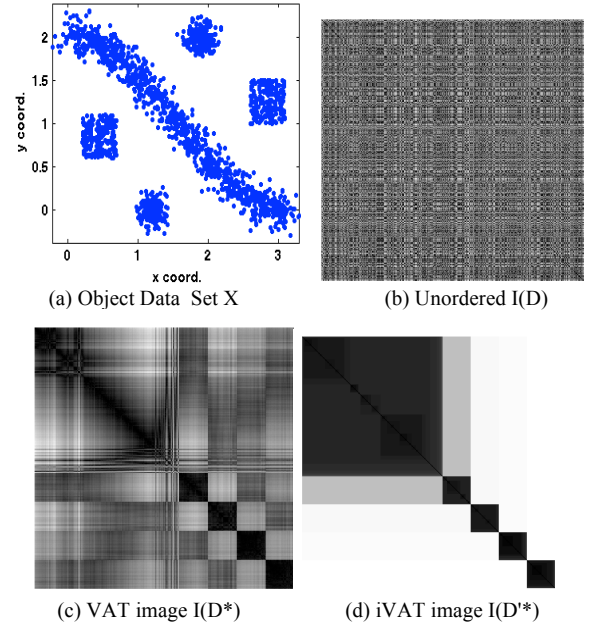


Figure 2. VAT/iVAT images of Boxes and Stripe

to all three groups, '4' belongs to groups 1 and 2, '5' belongs to group 1, while '1' and '2' are in just group 3. Thus, the joint membership of an individual in various communities is a function of the threshold λ . Zhang et al. do not specify the range of λ , but it must be $0 < \lambda < 1$, for otherwise the bounds of this conversion procedure would be $[1]_{c \times n}$ at $\lambda \leq 0$ and $[0]_{c \times n}$ at $\lambda \geq 1$. There are (infinitely) many candidate partitions available from this procedure because we can apply this process to candidates generated by FCM at each $c = 2, 3, \dots, n-1$; and within each c , for any $0 < \lambda < 1$. Zhang et al. choose a "best" possibilistic U by maximizing their version of the fuzzy modularity index, which is defined as follows:

$$V_k = \{i \mid u_{ik} > \lambda; 1 \leq i \leq c\}; \quad k = 1, \dots, n, \quad (8a)$$

$$S_z(V_k, V_k) = \sum_{i, j \in V_k} \left(\frac{u_{ik} + u_{jk}}{2} \right) w_{ij} \quad ; \quad (8b)$$

$$S_z(V_k, V) = S_z(V_k, V_k) + \sum_{\substack{i \in V_k \\ j \in V - V_k}} \left(\frac{u_{ik} + (1 - u_{jk})}{2} \right) w_{ij} \quad ; \quad (8c)$$

$$Q_z = \sum_{k=1}^c \left[\frac{S_z(V_k, V_k)}{S(V, V)} - \left(\frac{S_z(V_k, V)}{S(V, V)} \right)^2 \right]. \quad (8d)$$

The values $\{u_{ik}\}$ appearing in (8b) and (8c) are from the fuzzy c -partition before possibilistic conversion. Clearly, (8b) and (8c) reduce to $S(V_k, V_k)$ and $S(V, V_k)$, respectively, for crisp partitions (assuming $0 < \lambda < 1$). Hence, $Q_z = Q_h$ for crisp partitions. However, we believe that Q_z has theoretical problems, which we outline in [15]. Here, we are content to compare the analysis of two real data sets using the indices Q_z and Q_g .

IV. VISUAL CLUSTER TENDENCY ASSESSMENT

Let D be a set of square or rectangular dissimilarity data. The idea of visually analyzing the rows and/or columns of D to reveal structural relationships between individuals associated with D began with Loua [18] in 1873. The first *reordered dissimilarity image* (RDI) of a square data matrix appears in Czekanoski [19]. The methods for constructing and using RDIs in various applications have subsequently grown almost without bounds. Wilkinson and Friendly [20] survey contemporary methods using this idea in bioinformatics, where the reordered image is called a "cluster heat map." These authors state that this method has appeared in more than 4,000 papers in the last decade. Liiv [21] gives a very useful survey of seriation methods for social network analysis.

The *visual assessment of tendency* (VAT, [22]) model reorders symmetric, square D to D^* using the indices of a minimal spanning tree on D , and then displays a heatmap image $I(D^*)$ of D^* (often a gray-scale image). The basic rationale for VAT is that if an object tends to cluster with other objects, then it should also be part of a submatrix of "similarly small" values corresponding to those objects. These submatrices are seen as dark blocks along the diagonal of the VAT image $I(D^*)$. Contrast can be improved by setting the diagonal to the minimum of the off-diagonal values. Zhang et al. [23] discuss the use of VAT in a product called RoleVAT, a role engineering tool for role based access control. *Improved VAT* (iVAT, [24]) transforms D to D' using geodesic distances to replace the input distances, followed by VAT reordering of D' to D^* . We will also discuss a new adaptation of specVAT, a member of the VAT family related to spectral clustering [25].

Figure 2 illustrates VAT / iVAT using the 2D object data set X called *Boxes and Stripe* in [24]. View 2(a) scatterplots X , which is converted to the symmetric matrix D using the Euclidean norm, $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$. Figure 2(b) is the image $I(D)$. Figure 2(c) is the VAT reordered image $I(D^*)$, and Figure 2(d) is the iVAT reordered image $I(D^*)$. Most observers would agree that there are $c = 5$ pretty distinct clusters in *Boxes and Stripe*. This substructure is not evident in $I(D)$. The VAT image $I(D^*)$ does a better job at highlighting the structure, but the upper left block along the diagonal, which corresponds to the stripe cluster in X , is not so clear. More generally, the image in view 2(c) lacks clarity—there is not much contrast between the on and off diagonal blocks. Replacing Euclidean distances in D by geodesic distances in D' prior to VAT reordering with iVAT recursion renders the substructure in X quite nicely, as seen in Figure 2(d).

However, we will see that iVAT doesn't work as well for the Karate Club data set because social network data in the form of the graph $G = (V, E, W)$ don't respond well to iVAT reordering. In brief, the conversion of G to a distance matrix D is fraught with problems—we describe this problem more in detail in the next section. We now turn to describing the Karate Club Data and analyzing the use of various visual

clustering tools, including iVAT and specVAT, and fuzzy modularity.

V. NUMERICAL EXAMPLE: KARATE CLUB DATA

The Karate Club data is an undirected graph $G_k = (V, E, W)$ with 34 vertices that show links between the 34 members of a university karate club collected by Zachary in 1977 [4]. Edge weight w_{ij} indicates the relative strength of the association between individuals i and j (number of situations in and outside the club in which interactions occurred). The maximum value in W is 7 for the edge between members 26 and 32.

The Karate Club data is a favorite amongst social network-ists, because the evolution of the relationship between pairs of members in the Karate club—which was known and recorded by Zachary—provides a sort of "ground truth" for various social network analyses.

Zachary used these data and an information flow model of network conflict resolution to explain the split-up of this group into two factions (the squares and circles) following disputes among the members. The principals in the split were the karate instructor (vertex 1) and the president of the club (vertex 34).

Figure 3 shows Zachary's karate club network as depicted by Newman and Girvan [14]. Square nodes represent the instructor's faction and circular nodes the president's faction. The original belief was that this network should decompose well into two clusters as shown in Figure 3, because its members did bifurcate, following either the president or the instructor. However, various discussions of this data in the literature disagree.

Figure 4(a) is the image of the matrix $D=[7]-W$ (the diagonal of D is also set to 0 following this transformation) where W is the matrix of edge weights for the graph G_k . Figure 4(b) is the iVAT image of D . View 4(b) suggests that the Karate club has three pretty tight clusters (the red blocks), overlain by a weaker and larger orange block. The orange block is inside an even larger yellow block. Five individuals at the bottom and one at the top are isolated pixels in the iVAT image, indicating

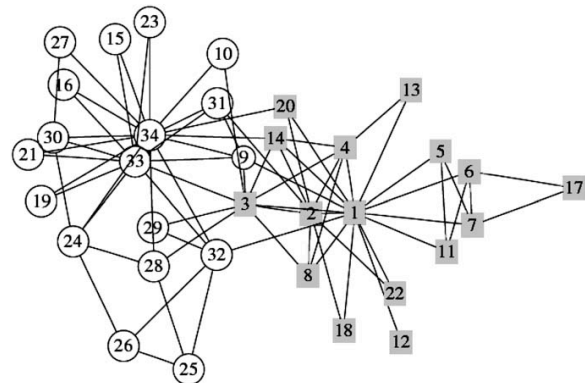


Figure 3. Karate club social network [14]

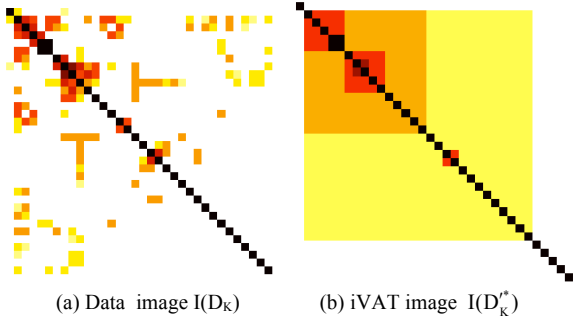


Figure 4. iVAT image of Karate club data

non-association with the other 28 members of the club. However, the implications of this representation of G_k are wrong. Essentially, the transformation of W to the distance matrix D considers the edges with 0-valued weight $w_{ij} = 0$ to have finite distance $d_{ij} = 7$, where, we argue, that these distances should be infinite (i.e., the absence of a path between i and j).

We examined other transformations of W into D , such as artificially increasing the distances corresponding to zero-weight edges, but did not achieve a pleasing result. Hence, we now turn to a spectral method for visualizing cluster tendency.

To alleviate the problem that we see with using VAT / iVAT with these data, we will use specVAT [25], which displays a dissimilarity image of the spectral components of a normalized weight matrix. specVAT first computes the top k eigenvectors of the eigenvalue problem

$$Lx = \lambda x, \quad (9)$$

where $L = M^{-1/2}WM^{-1/2}$ and M is the $n \times n$ diagonal matrix with the vector $\mathbf{m} = W\mathbf{1}_n$ on the diagonal. Then each of the k \mathbf{v} 's are normalized to the unit hypersphere by l_2 normalization. Finally, the distance matrix D is computed by $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$. VAT (and iVAT) can then applied to D to visualize the clustering tendency of W .

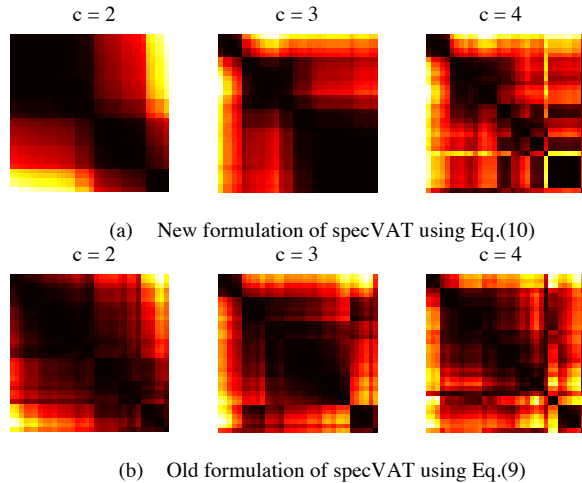


Figure 5. specVAT images of Karate Club data

We generalize specVAT by computing the eigenvectors \mathbf{x} with the generalized eigenvalue problem

$$Wx = \lambda Mx. \quad (10)$$

It can easily be shown that (10) is equivalent to (9), but for most eigensolvers, (10) is a more stable instantiation because small elements of \mathbf{m} do not induce numerical stability issues.

Figure 5 compares the use of (9) and (10) in creating the specVAT images using $c = 2, 3$, and 4 eigenvectors for the Karate Club data. It is somewhat clear from the specVAT images in view 5(a) that the Karate Club data have 3 clusters—there seem to be three dark blocks on the $c = 2$ and $c = 3$ images. At $c = 4$, the image begins to break down which indicates that there shouldn't be cluster structure at $c = 4$. In the specVAT images, shown in view 5(b), the cluster structure is less apparent. Interestingly, Figure 5 shows that although (9) and (10) are mathematically equivalent, the instantiation in the eigensolver can prove to drastically alter the results.

Although the specVAT images in Figure 5(a) seem to suggest 3 clusters, we wanted to test how the iVAT geodesic distance transform could be applied to improve the visualization. We applied the transformation to the distance matrix computed by specVAT. We call the formulation of specVAT using (10) with iVAT, specieVAT for *spectral improved eigensolver* VAT. These images are shown in Figure 6. It is now very clear, by the images in view 6(a), that specieVAT reinforces the popular viewpoint that the Karate Club data have 3 clusters. However, the views in 6(b), made by applying the iVAT transformation to the old formulation of specVAT, do not seem to suggest any visually-pleasing (relative to the precedent) cluster structure in the data set. We see this as further evidence that specVAT using (10) is superior to the original formulation—albeit mathematically the same. Furthermore, this shows that using the iVAT transformation with the new formulation of specVAT improves the visualization.

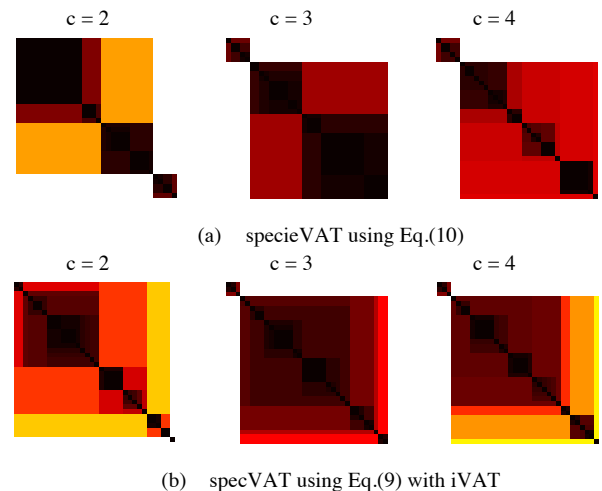


Figure 6. specieVAT and specVAT + iVAT images of Karate Club data




Figure 7. Crisp HCM 3-partition of X_k [12]




Figure 8. "Fuzzy" communities in G_k [12]

Now we turn to clusters in the vertex set of G_k . According to Zhang et al. [12], hard c -means applied to a vector representation X_k of G_k —which is computed by normalizing the top k eigenvectors of (10)—identifies the following 3 crisp clusters:

$$\begin{aligned} A &= \{5, 6, 7, 11, 17\} \\ B &= \{1, 2, 3, 4, 8, 12, 13, 14, 18, 20, 22\} \\ C &= \{9, 10, 15, 16, 19, 21, 23-34\} \end{aligned}$$

Figure 7 shows this crisp 3-partition of the Karate club data. Zhang et al. chose $c = 3$ from candidate partitions for $c = 2, 3, 4$ and 5 by maximizing the modularity function at (4). Then they ran FCM on X_k for the same values of c , and again selected $c = 3$ based on maximizing their fuzzy modularity index at (8d). Conversion of the fuzzy 3-partition of the vertices in G_k with $\lambda = 0.25$ led them to the possibilistic partition $U_{0.25} = \{A^* \cup B^* \cup C^*\}$ in Figure 8, which they called "fuzzy communities" in Karate Club. This diagram shows that 30 of the 34 individuals are assigned (crisply) in just one of the three clusters, while the four individuals numbered 9, 10, 31 and 1 (encircled in orange) are linked across two clusters. Zhang et al. did not specify which two clusters shared these members, but we infer from the positions in their diagram that members 9, 10 and 31 had full membership in B^* and C^* while the instructor (vertex 1) had full membership in B^* and A^* .

To show the comparison of Q_h , Q_z , and Q_g we implemented the modification of the MULTICUT spectral clustering algorithm described in [26], applying FCM to the resulting spectral features extracted from the graph G . To compute Q_h we hardened the resulting fuzzy partition and applied equation (6). For Q_z , we will show values for $\lambda = 0.1$ and 0.25 . The fuzzy modularity at (7) was computed directly on the fuzzy partition U . For all experiments, we set the FCM fuzzifier $m = 2$. FCM was initialized by randomly selecting c vectors from the data as initial cluster centers. We then ran the clustering algorithm for $c = 2, 3, \dots, 10$ (for example). The indices were calculated for each c -partition and the maximum determined the chosen value of c . We performed this experiment 50 times. The plots here shown the mean value of each modularity index over the 50 runs at each c . The table indicates the number of times that an index maximized at each c over the 50 experiments.

Figure 9 and Table I show the normalized results (max set to 1) for the three modularity indices on the Karate Club data set. When examining this plot, remember that the overall value is not as important as the value of c at which each line peaks. The plot shows that Q_g is the only index that clearly prefers $c = 3$. But the table clearly show that both Q_g and $Q_z(0.25)$ prefer $c = 3$ most of the time. The indices Q_h and $Q_z(0.1)$ prefer $c = 4$. The plot and table show that Q_z is more uncertain of its choice of c . Does this mean that $Q_z(0.1)$ is inferior? No. The "true" cluster structure of these data is unknown—many have only postulated, based on empirical and anecdotal evidence, that $c = 3$. The generalized modularity backs up this

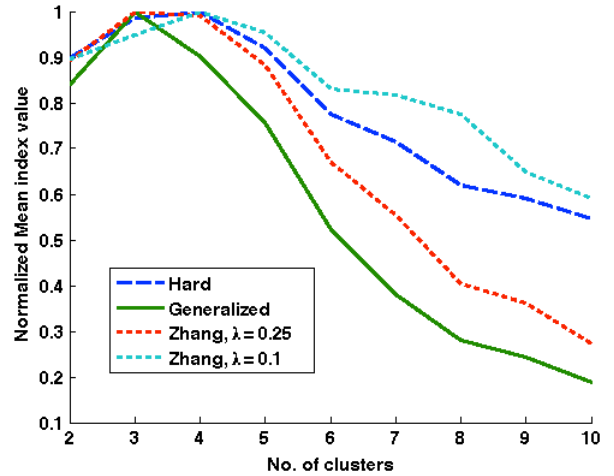


Figure 9. Mean values of three modularity indices for 50 runs of the Karate Club data set over $c = 2$ to 10 . Tabulated values indicate number of times each index maximized at each c .

TABLE I. NUMBER OF TIMES EACH MODULARITY INDEX MAXIMIZED AT c OVER 100 RUNS.

	2	3	4	5	6	7	8	9	10
Q_h	0	8	91	1	0	0	0	0	0
Q_g	0	93	7	0	0	0	0	0	0
$Q_z(0.25)$	2	60	37	1	0	0	0	0	0
$Q_z(0.1)$	2	7	70	15	2	2	2	0	0

belief; however, we would say that number of clusters in this data set is uncertain. If we simply add up the columns of Table I, $c = 4$ is the clear choice. But if we examine the plot in Fig. 9, Q_g has the sharpest peak, occurring at $c = 3$. This is a prime example of the cluster validity conundrum—which partition do I choose and which validity measure do I trust? There is no perfect answer to this question, but this experiment does point to an underlying issue in Q_z ; namely, how does one choose λ ? In this case, it would be easy, in hindsight, to say that one should choose $\lambda = 0.25$ —this value of λ chooses the preferred $c = 3$. However, what works for this data set may not work for every data set; hence, we prefer the parameter-free generalization Q_g .

VI. CONCLUSIONS

Our generalization of Newman’s modularity function at (7) represents a step-forward in the problem of finding fuzzy communities in graph data. First, it is a direct generalization of the underlying theory of graph modularity: namely, providing a measure of the probability of a graph structure relative to a null hypothesis. For more on this discussion, please refer to [15]. Second, it is a parameter-free instantiation of fuzzy modularity: the only of its kind to date. Last, we showed that on the well-studied Karate Club data set that our fuzzy modularity index performs comparably to Newman’s crisp modularity function and to Zhang et. al’s fuzzy index.

The second point of progress in this paper is the reformulation of specVAT into the new algorithm, specieVAT, for determining the number of clusters in graph-based data. First, we showed that specVAT can be reformulated as a generalized eigenvalue problem at (10). This formulation is mathematically equivalent to the original specVAT, but is numerically more stable for most, if not all, eigensolvers. Second, we showed how the iVAT distance transform could be applied to specVAT to improve the tendency assessment visualization.

In the future, we will examine how our fuzzy modularity at (7) can be directly maximized. Initial work in this direction indicates that (7) can be posed as a generalized eigenvalue problem much like (10). Our current efforts are focused on establishing this maximization problem within the existing spectral clustering framework and on stabilizing the numerical issues involved in reaching the solution. Second, we will examine how fuzzy modularity can be generalized as a cluster validity index (or quality function) for asymmetric adjacency matrices or directed graphs. Finally, we will do a comprehensive comparison of various validity indices and clustering algorithms on a wide assortment of graph-based and social network data. As always, we believe that there is no free lunch in the clustering game and that the best attacks at these problems will involve a menagerie of clustering tools.

REFERENCES

[1] S.F. Sampson, A novitiate in a period of change: an experimental and case study of social relationships. Unpublished Doctoral Dissertation. 1968.

[2] R. Breiger, S. Boorman, P. Arabie, "An algorithm for clustering relational data with applications to social network analysis and comparison with multidimensional scaling," *Journal of Mathematical Psychology*, vol. 12, 1975, pp. 328-383.

[3] S. Wasserman and K. Faust, *Social Network Analysis*. Cambridge: Cambridge University Press, 1994.

[4] W. Zachary, "An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*," vol. 33, 1977, pp. 452-473.

[5] L. A. Zadeh, "Similarity relations and fuzzy orderings," *Information Sciences*, vol. 3, 1971, pp. 177-200.

[6] J. M. Blin, "Fuzzy relations in group decision theory," *J. Cybernetics*, 1974, pp. 17-22.

[7] J. C. Bezdek, B. Spillman, B. and R. Spillman, "A Fuzzy Relation Space for Group Decision Theory", *Fuzzy Sets and Systems*, vol. 1, 1978, pp. 255-268.

[8] J. C. Bezdek, B. Spillman and R. Spillman, "Fuzzy relation spaces for group decision theory: An application," *Fuzzy Sets and Systems*, vol. 2, 1979, pp. 5-14.

[9] B. Spillman, J. Bezdek and R. Spillman, "Development of an Instrument for the Dynamic Measurement of Consensus," *Comm. Mono.*, vol. 46, 1979, pp. 1-12.

[10] S. Fortunato, "Community detection in graphs." *Phys. Rep.*, vol. 486, 2010, pp. 75-174.

[11] J. Reichardt and S. Bornholdt, "Detecting fuzzy community structures in complex networks with a Potts model," *Stat. Mech*, 2004.

[12] S. Zhang, R.S. Wang and X.S. Zhang, "Identification of overlapping community structure in complex networks using fuzzy c-means clustering," *Stat. Mechanics and its Appl.*, vol. 374(1), 2007, pp. 483-490.

[13] J.C. Bezdek, J.M. Keller, R. Krishnapuram, and N.R. Pal, *Fuzzy models and algorithms for Pattern Recognition and Image Processing*, Springer: NY, 1999.

[14] M. E. J. Newman and M. Girvan, *Phys. Rev. E* 69(2), 2004.

[15] T. C. Havens, J. C. Bezdek, C. Leckie, K. Ramamohanarao, and M. Palaniswami, "A soft modularity function for detecting fuzzy communities in social networks," *IEEE Trans. Fuzzy Systems*, 2013, doi:10.1109/TFUZZ.2013.2245145.

[16] U. Brandes, D. Delling, M. Gaertler, R. Gorke, M. Hofer, Z. Nikoloski and D. Wagner, "On Modularity Clustering," *IEEE TKDE*, vol. 20(2), 2008, pp. 172-188.

[17] J. Liu, "Fuzzy modularity and fuzzy community structure in networks," *Eur. Phys. J. B*, vol. 77, 2010, pp. 547-557.

[18] T. Loua Atlas statistique de la population de Paris, J. Dejeu, 1873.

[19] J. Czekanowski, "Zur differentialdiagnose der neandertalgruppe," *Korrespondenzblatt der Deutschen Gesellschaft fr Anthropologie, Ethnologie und Urgeschichte*, vol. 40, 1909, pp. 44-47.

[20] L. Wilkinson and M. Friendly, "The history of the cluster heat map," *The American Statistician*, vol. 63(2), 2009, pp. 179-184.

[21] I. Liiv, "Seriation and Matrix Reordering Methods: An Historical Overview," *Stat. Anal. and Data Mining*, vol. 3(2), 2010, pp. 70-91.

[22] J. C. Bezdek and R. J. Hathaway, "VAT: A tool for visual assessment of (cluster) tendency," in *Proc. IJCNN*, 2002, pp. 2225-2230.

[23] D. Zhang, K. Ramamohanarao, S. Versteg and R. Zhang, "RoleVAT: Visual Assessment of Practical Need for Role Based Access Control," *Proc. IEEE Conf. on Security Apps*, 2009, pp. 13-22.

[24] T. C. Havens and J. C. Bezdek, "An Efficient Formulation of the Improved Visual Assessment of Cluster Tendency (iVAT) Algorithm," *IEEE TKDE*, vol. 24(5), 2012, pp. 813-822.

[25] L. Wang, X. Geng, J.C. Bezdek, C. Leckie and K. Ramamohanarao, "Enhanced Visual Analysis for Cluster Tendency Assessment and Data Partitioning," *IEEE TKDE*, vol. 22(10), 2010, pp. 1401-1414.

[26] D. Verma and M. Meila, "A comparison of spectral clustering algorithms," *UW CSE Technical Report*, 03-05-01, 2003.