



GANC: Greedy agglomerative normalized cut for graph clustering

Seyed Salim Tabatabaei*, Mark Coates, Michael Rabbat

Department of Electrical and Computer Engineering, McGill University, McConnell Engineering Building, 3480 University Street, Montreal, Quebec, Canada H3A 2A7

ARTICLE INFO

Article history:

Received 26 November 2010

Received in revised form

30 April 2011

Accepted 13 June 2011

Available online 18 July 2011

Keywords:

Graph clustering

Normalized cut

Model order selection

Large scale graphs

ABSTRACT

This paper describes a graph clustering algorithm that aims to minimize the normalized cut criterion and has a model order selection procedure. The performance of the proposed algorithm is comparable to spectral approaches in terms of minimizing normalized cut. However unlike spectral approaches, the proposed algorithm scales to graphs with millions of nodes and edges. The algorithm consists of three components that are processed sequentially: a greedy agglomerative hierarchical clustering procedure, model order selection, and a local refinement.

For a graph of n nodes and $O(n)$ edges, the computational complexity of the algorithm is $O(n \log^2 n)$, a major improvement over the $O(n^3)$ complexity of spectral methods. Experiments are performed on real and synthetic networks to demonstrate the scalability of the proposed approach, the effectiveness of the model order selection procedure, and the performance of the proposed algorithm in terms of minimizing the normalized cut metric.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Clustering or partitioning of nodes in networks or graphs is an important task that has many applications in a diverse range of fields. It has been used for many years to study social networks [1] and continues to be employed in the field of sociology to explore social interactions. More recently it has been employed in the study of biochemical networks [2,3], biological neural networks [4,5], and transport and communication networks.

There are three components to the clustering task: (i) choosing the number of clusters; (ii) selecting a criterion that measures the merit of each candidate cluster allocation; and (iii) identifying an algorithm that searches for the optimal clustering. Some performance criteria can be used both to select the number of clusters and choose the clustering, e.g., modularity [6] or information-theoretic criteria based on the minimum description length [7].

There is no universally accepted performance criterion and indeed the most appropriate criterion can vary depending on the application domain and the goal of the clustering. Modularity, for example, focuses on network structure and places primary value on direct connections between nodes. On the other hand, clustering algorithms based on Markov random walks [7,8] also value indirect connections and network flow. In this paper we select the normalized cut criterion [9], which simultaneously encourages intra-cluster similarity while penalizing inter-cluster similarities.

Methods based on this criterion have been employed successfully in a wide range of applications [5,8–12]. The normalized cut criterion is related to the conductance of the underlying graph [13]. Furthermore an implicit duality between normalized cut and normalized association exists; the former encourages clusters that are less connected to each other and the latter encourages clusters whose nodes are well-connected.

When partitioning a graph into clusters, the cut associated with cluster i is the sum of the weights of the edges between nodes in i and nodes in other clusters. Intuitively, minimizing the maximum cut (or minimizing the average cut) identifies clusters that have the weakest ties between them. The problem is that the minimum cut will often be achieved by identifying many very small clusters, which provides little insight into the underlying structure of the graph. The normalized cut metric was introduced by Shi and Malik in [9] to address this shortcoming. It normalizes each cut by dividing it by the total weight of the associated cluster. This has the effect of penalizing very small clusters because they generally have low total weight.

Minimizing normalized cut is an NP-complete problem [9]. Shi and Malik illustrated that the minimization could be relaxed to form a generalized eigenvalue system, whose (discretized) solution corresponds to the minimum normalized cut. This has led to the development of spectral clustering; it involves calculating eigenvectors of the identified system. In order to identify a partitioning of n nodes to k clusters, some techniques perform recursive bipartitioning [9], and thus require the repeated identification of two eigenvectors. Other approaches strive to identify k clusters directly by calculating the k smallest eigenvectors of the underlying graph Laplacian. One of the concerns about these

* Corresponding author. Tel.: +1 514 677 0056; fax: +1 514 398 4470.

E-mail addresses: seyed.s.tabatabaei@mail.mcgill.ca (S.S. Tabatabaei), mark.coates@mcgill.ca (M. Coates), michael.rabbat@mcgill.ca (M. Rabbat).

methods is computational complexity; even when employing fast eigendecomposition methods such as the Lanczos iterative technique, complexity grows rapidly with n .

In this paper, we propose an agglomerative clustering algorithm that strives to minimize the normalized cut (or equivalently, maximize the normalized association). It is a fast, scalable algorithm, with almost linear complexity in the number of nodes for relatively sparse graphs. Performance evaluation using a range of benchmark and observed graphs indicates that the algorithm identifies partitionings that have average normalized association metrics as large as those of the partitions identified by the spectral clustering techniques. We also propose a method for identifying important partitioning scales which can be used to automatically select the number of clusters. To the best of our knowledge, this is the first model order selection method applicable to the normalized association maximization that is scalable. Through our experiments, we will show how effective our model order selection criterion is for both synthetic and real networks.

The rest of the paper is organized as follows. In Section 2 the major approaches for minimization of normalized cut and model order selection are reviewed. The problem formulation of normalized association is re-stated in Section 3. Our proposed algorithm is detailed in Section 4. The proposed algorithm is then compared to the state-of-the-art clustering algorithms in Section 5. The concluding remarks are provided in Section 6.

2. Related work

The identification of clusters in graphs and networks has received significant attention. In our review we focus on a representative set of algorithms that minimize the normalized cut with or without spectral decomposition. We also review the existing methods for selecting the number of clusters.

2.1. Optimization of normalized cut

Spectral partitioning of graphs was first proposed by Donath and Hoffman in 1970s [14]. Interest in the techniques was renewed in the 1990s when Pothen et al. described an algorithm for bipartitioning using the Fiedler vector [15]. Hendrickson et al. and Karypis et al. contributed with multilevel algorithms for more efficient spectral partitioning [16,17].

The normalized cut metric was introduced by Shi and Malik in [9]. They demonstrated how the bipartitioning task, with the objective of minimizing the normalized cut, could be relaxed to construct a generalized eigenvalue problem, and was thus related to spectral partitioning. The eigenvector corresponding to the second-smallest eigenvalue of the graph Laplacian identifies the bipartitioning (the real values in the eigenvector must be mapped to two discrete values for partitioning). This established the connection between minimization of normalized cut and spectral partitioning. Shi and Malik proposed a recursive bipartitioning scheme in order to partition a graph into k clusters. In [18], Meila and Shi proposed an algorithm that calculates k eigenvectors (thereby associating k real values with each node in the graph) and then uses a clustering algorithm, such as k -means, to do the partitioning in \mathfrak{R}^k . Ng et al. observed in [19] that the algorithm in [18] is susceptible to failure when there is substantial variation in the degree of connectivity between clusters. They proposed an alternative algorithm that uses a different normalization in both the eigenvalue problem and the construction of feature vectors prior to the application of k -means.

All of these algorithms involve a computationally expensive eigendecomposition. To address the computational difficulties for large graphs, Fowlkes et al. described a procedure that uses the

Nyström method to reduce the complexity of the eigenvalue problem [10]. This does significantly reduce the computational overhead, but it is not enough to make the eigendecomposition algorithms scalable to large graphs. Yan et al. recently described an algorithm for fast approximate spectral clustering [20], but the focus is not on clustering for graphs (rather it addresses real-valued feature vectors).

Dhillon et al. introduced a much faster algorithm for minimizing normalized cut in [11]; the graph is first greedily coarsened, then the coarsened graph is partitioned using a region growing procedure [17] and finally weighted kernel k -means clustering is applied to each partition to refine the clustering.

Other methods also exist that strive to optimize related cost functions; for example Sharon et al. proposed a scalable hierarchical clustering using *ratio association* (the normalization term is the number of nodes in clusters) [21]. Other examples include the work by Ding et al. [22], Sharon et al. [23], Akselrod-Ballin et al. [24], and Corso et al. [25] that use the sum of internal weights of clusters as the normalization term.

2.2. Model order selection

The problem of selecting the number of clusters is almost as old as clustering itself. In the following discussion, k^* denotes the number of clusters that a given approach identifies as the true number of clusters.

Numerous authors have used some notion of quality of clustering, usually based on some definition of the inter- and intra-cluster distances, to identify the number of clusters. Let $F(k)$ denote this notion for a given partitioning to k clusters. $F(k)$ is then examined for various values of k . Partitionings for different values of k can be obtained from a hierarchical clustering, or alternatively, by several flat partitionings.

If F is not a monotonic function of k , the approach is usually to take $k^* = \arg \max_k F(k)$ or $k^* = \arg \min_k F(k)$ (depending on the definition of F). The majority of the methods discussed in [26,27] are of this type; Calinski and Harabasz [28], C-index [29], Baker and Hubert [30], and cubic clustering criterion [31] are among the most effective ones [26]. Tibshirani et al. [32] define $F(k)$ as the *gap* that is the difference of the average pairwise distance of the data points of the clustering at k th level and the expected value of the same measure of some reference model. This is similar to modularity [6] (discussed later) in the sense that it compares the clustering results to a reference model. None of these methods are directly applicable to graph clustering algorithms; the calculations of the defined metrics require pairwise distances which are not immediately available from a graph representation. Possible distance metrics include the shortest path [33] or the diffusion distance [8]; however the shortest path is very sensitive to noise and the calculation of the diffusion distance requires eigendecomposition.

In the case where F is a monotonic function of k , the only extremal values of $F(k)$ correspond to trivial values of k . Hence, the value of $F(k)$ corresponding to two or more choices of k are examined to quantify the significance of a given level of a hierarchical clustering [34–36]. In order to assess the clustering at level k of the hierarchy, Gnanadesikan et al. [34] propose the fraction $F(k)/F(k-1)$, Arnold [37] uses the value of $F(k)-F(k-1)$, Krzanowski and Lai [35] employ the fraction $|F(k)-F(k-1)|/|F(k+1)-F(k)|$, and Pederson and Kulkarni [36] suggest using the fraction $2 \times F(k)/(F(k-1)+F(k+1))$. The latter two approaches are most similar to what we propose in Section 4 in the sense that they use the preceding and succeeding levels of the hierarchy to obtain the significance of the clustering at level k . However the approaches in [35,36] are potentially susceptible to noise due to the division; small perturbations in the weights could lead to dramatic changes in the selected model order.

Other approaches to selecting the number of clusters are based on the adoption of (semi-)parametric models for the structure of the graph. This allows the application of model selection techniques based on concepts such as the Bayesian Information Criterion (BIC) [38] and Akaike Information Criterion [39]. Despite the strong theoretical support for these methods, the adoption of a parametric model for the graph structure is undesirable. The models are often overly restrictive and do not adequately capture the properties of many real-life networks. An example is the requirement in [40,41] that the input data are normally distributed (after projection of the graph into a real space). The more general methods based on mixture models do not scale well to very large graphs; even the recent approaches have only been applied to graphs with a few thousand nodes [42–44].

Some heuristics are based on the sizes of the clusters that are merged at different levels of the clustering hierarchy [45,46]. The authors in [46] suggest that when two clusters with large number of nodes are merged, a significant amount of detail is lost; hence a point at which this occurs is probably a good place for a hierarchical clustering algorithm to stop. The authors of [45] propose a similar approach. The main drawback of these approaches is that only the granularity of the clusters are taken into account and the number and the weights of the edges are simply ignored.

A well-known and effective method of selecting the number of clusters is to examine the eigenvalues of the Laplacian of the graph that is to be clustered [19,47,48]. For a graph with isolated connected components the multiplicity of the eigenvalue zero is equal to the number of clusters. Any other graph with well-separated clusters can be considered as a perturbation of this ideal case. Matrix perturbation theory states that the stability of the eigenvectors of a matrix is proportional to the eigengap (the difference between two successive eigenvalues). Von Luxburg [48] suggests using $k^* = \arg \max_k (\lambda_k - \lambda_{k+1})$. A large eigengap at k^* is the case in which spectral algorithms using the Laplacian perform most successfully [19]. A more robust criterion is proposed by Zelnik-Manor and Perona [49] that uses eigenvectors instead. Despite the solid theoretical support behind these approaches, the requirement of eigendecomposition makes them impractical when large graphs are being clustered.

There have been attempts to identify the number of clusters using stability analysis. “Stability” has been defined differently by different authors, but they all based on the same intuition; with respect to some algorithm, a stable clustering is one that behaves relatively consistently in the presence of some controlled perturbation. The perturbation could be in terms of noise [50], sampling subsets from the input [51,52], or random projection of high dimensional data into a lower dimensional space [53]. These approaches require several runs of clustering for every value of k which makes them computationally expensive. Furthermore, Ben-David et al. [52] warn against using stability analysis in this context, and they suggest that this family of model order selection techniques is not suitable for selecting the number of clusters in general. The intuition behind their work is that when an underlying objective function, $F(k)$, has several local optima of relatively similar values, a clustering algorithm might be trapped by any of them, even when k is the true number of clusters. The difference in clustering solutions is interpreted as instability and the model order rejected, but the behavior is caused by the imperfect nature of the clustering algorithm.

Two more recently proposed methods to select the number of clusters are based on optimization of the quality metrics of modularity [6] and description length [7]. Both methods simultaneously address both clustering and the selection of the number of clusters; they strive to optimize the objective function over all possible partitionings. These methods are discussed in more detail in Section 5.

3. Problem formulation

Let $G = (V, E, w)$ be a weighted graph having $n = |V|$ nodes and $m = |E|$ edges.¹ We assume that edge weights $w(u, v) = w(v, u) \geq 0$ are non-negative and symmetric, that $w(u, v) = 0$ if $(u, v) \notin E$, and that for $(u, v) \in E$, the weight $w(u, v) \geq 0$ is indicative of the similarity between nodes u and v ; that is, the larger the weight, the more similar the nodes. We also allow for self-weights, $w(u, u) \geq 0$.

For a fixed number of clusters, k , we measure the quality of the partition via the normalized cut metric [9], defined as follows. For a node u , denote the degree of u by $d(u) = \sum_{v \in V} w(u, v)$, and for a subset $U \subset V$ of nodes, let $d(U) = \sum_{u \in U} d(u)$ denote the cumulative degree of the subset. Similarly, for two disjoint subsets of nodes, V_1 and V_2 , let $w(V_1, V_2) = \sum_{u \in V_1} \sum_{v \in V_2} w(u, v)$ denote the sum of the weights of edges with one end in V_1 and the other end in V_2 . Let C_k denote the partition of nodes to k clusters, $C_k = \{C_1, \dots, C_k\}$ where C_i is the subset of nodes affiliated to cluster i . The *normalized cut metric* is defined as

$$\text{NCut}(C_k) = \sum_{i=1}^k \frac{w(C_i, V \setminus C_i)}{d(C_i)}. \quad (1)$$

Minimizing the normalized cut can be interpreted as minimizing the similarity of nodes in different clusters, relative to the degree of each cluster. Alternatively, maximizing the intra-cluster similarity can be achieved by maximizing the *normalized association*, defined as

$$\text{NAssoc}(C_k) = \sum_{i=1}^k \frac{w(C_i, C_i)}{d(C_i)}. \quad (2)$$

Moreover, maximizing normalized association is equivalent to minimizing normalized cut since $w(C_i, C_i) + w(C_i, V \setminus C_i) = d(C_i)$, and hence $\text{NAssoc}(C_k) = k - \text{NCut}(C_k)$. For the sake of readability, we adopt the maximization of normalized association as our goal in developing a clustering procedure.

4. GANC: greedy agglomerative normalized cut

In this section, we describe a greedy algorithm for building an agglomerative clustering on a graph. Although we do not make guarantees about its accuracy, the algorithm is fast on sparse graphs and yields excellent performance on a variety of examples, as illustrated in Section 5. Moreover, GANC has a model order selection criterion and does not have to be provided with the number of clusters a priori. The algorithm consists of three steps: agglomerative clustering, model order selection, and refinement.

4.1. Agglomerative clustering

We propose greedy maximization of normalized association via an agglomerative hierarchical clustering. In the following discussion, note that the number of clusters at stage k of the hierarchy is k . At stage k of the hierarchy, using the given partition, C_k , two clusters are merged to form a new partition, C_{k-1} . The two clusters to be merged are chosen greedily so that normalized association of stage $k-1$ is maximized.²

¹ In this work, we assume we are given the graph on which we wish to perform clustering. We do not address the problem of forming a graph from data, which arises when one applies graph clustering methods to general data sets; see, e.g., [48,54,55].

² We note that Shi and Malik allude to a similar greedy merging algorithm in [9], however they also suggest first projecting each node into \mathbb{R}^k using the first k eigenvectors of the graph Laplacian, and running an algorithm such as k -means to obtain an initial clustering. Our approach requires no eigendecomposition and performs no such initialization. We also note that, although Shi and Malik mention

Initially, C_n is a function that maps nodes to unique clusters, $1, \dots, n$. The degree of every node u , $d(u)$, is calculated. Furthermore, for every edge, $(u, v) \in E$, the improvement in normalized association by its contraction to obtain C_{n-1} is stored in $\Delta(u, v) = 2w(u, v)/(d(u) + d(v))$, so that a matrix of the improvements is constructed. Here we assume that initially there are no self-loops, but if there are, the value of improvements are computed by (4) presented shortly. Throughout the agglomerative clustering procedure, $d(\cdot)$, $w(\cdot, \cdot)$, and $\Delta(\cdot, \cdot)$ are updated at each stage as described below.

At each iteration, the pair $(u^*, v^*) = \operatorname{argmax}_{(u, v)} \Delta(u, v)$ is selected for merging to create a larger cluster, uv^* . The degree is computed for the newly constructed cluster, $d(uv^*) = d(u^*) + d(v^*)$. The weights and the improvement matrices are updated by removing the rows and columns corresponding to u^* and v^* , and inserting a new row and column corresponding to uv^* (rows and columns are not removed or added in our implementation, but this is the practical effect). The weight matrix is updated as follows:

$$w(uv^*, x) = w(x, uv^*) = w(u^*, x) + w(v^*, x), \tag{3}$$

and the improvement matrix update is

$$\Delta(uv^*, x) = \frac{w(uv^*, uv^*) + w(x, x) + 2w(uv^*, x)}{d(uv^*) + d(x)} - \frac{w(uv^*, uv^*)}{d(uv^*)} - \frac{w(x, x)}{d(x)}, \tag{4}$$

for all the clusters, x , adjacent to either u^* or v^* . The self-weights are also calculated by:

$$w(uv^*, uv^*) = w(u^*, u^*) + w(v^*, v^*) + 2w(u^*, v^*). \tag{5}$$

For all pairs of nodes (u, v) not adjacent to u^* and v^* , the weights $w(u, v)$ and improvements $\Delta(u, v)$ remain unchanged. The above sequence of steps is repeated $n-1$ times to form the clustering hierarchy.

4.2. Model order selection

Many of the clustering algorithms require the number of clusters to be provided to the algorithm a priori. However such information is often not available in practical situations making the decision about the number of clusters an issue in itself.

It is worth noting that the stage number (k) that maximizes $\text{NAssoc}(C_k)$ does not necessarily correspond to a meaningful number of clusters.³ Here we propose a simple but effective approach to model order selection. Let $C_k^* = \operatorname{argmax}_{C_k} \text{NAssoc}(C_k)$ denote the partition that maximizes normalized association over all partitions of V into k clusters. To carry out model order selection, we examine the curvature⁴ of $\text{NAssoc}(C_k^*)$ which we define as

$$\begin{aligned} \text{Curv}(k) &= (\text{NAssoc}(C_k^*) - \text{NAssoc}(C_{k-1}^*)) - (\text{NAssoc}(C_{k+1}^*) - \text{NAssoc}(C_k^*)) \\ &= 2\text{NAssoc}(C_k^*) - \text{NAssoc}(C_{k-1}^*) - \text{NAssoc}(C_{k+1}^*). \end{aligned} \tag{6}$$

The first term of the above addition is the improvement of normalized association moving from stage k to $k-1$ of the hierarchy and the second term is the improvement moving from stage $k+1$ to k .

(footnote continued)

having experimented with greedy merging, results for this approach have not been reported in the literature.

³ To see this, consider an unweighted graph that consists of two isolated chains, each having four nodes. The true number of clusters is trivially 2 resulting in $\text{NAssoc}(C_2) = 2$; however if one groups the adjacent nodes to obtain four groups of node pairs, the resulting value of normalized association would be $\text{NAssoc}(C_4) = 2.66$.

⁴ The reason we call this metric the curvature is its similarity to the central approximation of the second order derivative which is defined for a continuous function, $f(\cdot)$ as $f''(x) \approx \frac{\partial^2 f(x)}{\partial x^2} = [f(x+h) - 2f(x) + f(x-h)]/h^2$. By substituting $h=1$, we get the negative of our curvature equation (the negation is to have positive peaks).

The function $\text{Curv}(k)$ captures the notion that a particular number of clusters, k , identifies meaningful structural similarities embodied in the graph if it provides a normalized association which is significantly larger than the best partition with one additional cluster ($k+1$) and little can be gained by reducing the number of clusters to $k-1$.

In practice we do not have access to the optimal partitions, so we cannot evaluate the exact value of the curvature function. Instead, we approximate the curvature by using the normalized association values for the partitions returned by the agglomerative step of our algorithm.

Note that the model order selection step of the algorithm could be used as a model order selection step for any other algorithm that maximizes the normalized association and generates a clustering hierarchy. Furthermore, this step of the algorithm can be considered optional if there is a prior knowledge about the true number of clusters.

4.3. Refinement

Greedy algorithms can get trapped by local optima. This is also the case with the agglomerative step of our algorithm, especially when the clusters are not clearly separated. After selecting the number of clusters either using the model order selection rule described previously or prior knowledge, a refinement step is invoked in order to improve the initial clustering results. The nodes are moved across the clusters to further improve the value of normalized association. A similar approach is taken in [56], but groups of nodes are moved from a cluster to another instead of individual nodes. If we try all possible moves, we end up performing an exhaustive search of dividing n nodes into k clusters. This defeats the purpose of developing the fast agglomerative clustering step. Instead, we look only at the boundary nodes, i.e., the nodes that have at least one neighbor in another cluster.

Initially the set of all the boundary nodes is identified. We use an $n \times k$ matrix to keep track of the neighborhood information of the boundary nodes. If B denotes this matrix, and I denotes an n -dimensional vector:

$$B(u, i) = \begin{cases} \sum_{v \in C_i} w(u, v) & \text{if } u \notin C_i \\ 0 & \text{otherwise} \end{cases}, \quad I(u) = \sum_{v \in C_i} w(u, v) \text{ if } u \in C_i. \tag{7}$$

For each of the boundary nodes, improvements in normalized association by moving from their current cluster to each of their neighbors are calculated. This improvement for node u moving from cluster i to cluster j is

$$\delta(u, i, j) = \frac{w(C_i, C_i) - 2I(u)}{d(C_i) - d(u)} + \frac{w(C_j, C_j) + 2B(u, j)}{d(C_j) + d(u)} - \left(\frac{w(C_i, C_i)}{d(C_i)} + \frac{w(C_j, C_j)}{d(C_j)} \right). \tag{8}$$

If no move leads to improvement, the node stays in the cluster to which it belongs. If one or more moves result in improvement, then the node is moved to the cluster that results in maximum improvement. If u is moved from C_i to C_j , the corresponding cluster degrees and associations are updated:

$$d(C_i)^{\text{new}} = d(C_i)^{\text{old}} - d(u), \tag{9}$$

$$d(C_j)^{\text{new}} = d(C_j)^{\text{old}} + d(u), \tag{10}$$

$$w(C_i, C_i)^{\text{new}} = w(C_i, C_i)^{\text{old}} - 2I(u, i)^{\text{old}}, \tag{11}$$

$$w(C_j, C_j)^{\text{new}} = w(C_j, C_j)^{\text{old}} + 2B(u, j)^{\text{old}}, \tag{12}$$

and for all the neighboring nodes of u denoted by v , entries of B and I are updated as follows:

$$\left. \begin{aligned} I(v)^{\text{new}} &= I(v)^{\text{old}} + w(u,v) \\ B(v,i)^{\text{new}} &= B(v,i)^{\text{old}} - w(u,v) \end{aligned} \right\} \text{ if } v \in C_j, \quad (13)$$

$$\left. \begin{aligned} I(v)^{\text{new}} &= I(v)^{\text{old}} - w(u,v) \\ B(v,j)^{\text{new}} &= B(v,j)^{\text{old}} + w(u,v) \end{aligned} \right\} \text{ if } v \in C_i, \quad (14)$$

$$\left. \begin{aligned} B(v,i)^{\text{new}} &= B(v,i)^{\text{old}} - w(u,v) \\ B(v,j)^{\text{new}} &= B(v,j)^{\text{old}} + w(u,v) \end{aligned} \right\} \text{ if } v \notin C_i \text{ and } v \notin C_j, \quad (15)$$

and finally

$$B(u,i)^{\text{new}} = I(u)^{\text{old}}, \quad (16)$$

$$I(u)^{\text{new}} = B(u,j)^{\text{old}}, \quad (17)$$

$$B(u,j)^{\text{new}} = 0. \quad (18)$$

While performing the updates (9)–(18), the set of boundary nodes is also updated. Whenever a node is moved from cluster i to cluster j , its neighbors in cluster i are added to the set of boundary nodes. For some node v neighboring u , moving u might result in $B(v,i)^{\text{new}}$ becoming zero in (15) for all i ; i.e., some nodes might be removed from the set of boundary nodes.

One pass through all the boundary nodes is considered a single refinement *iteration*. When an iteration causes no positive improvement, the refinement procedure is stopped. Alternatively, one could specify the maximum number of refinement iterations. Note that although we prohibit the refinement step from emptying any cluster, we have not observed such an attempt in our experiments.

Algorithm 1. GANC: greedy agglomerative normalized cut

GREEDY AGGLOMERATION

- 1: Build the initial Δ matrix, and
NAssoc(C_n) $\leftarrow \sum_{u \in V} w(u,u)/d(u)$
- 2: **for** $k = n-1$ to 1 **do**
- 3: $(i^*, j^*) \leftarrow \arg \max_{i,j} \Delta(i,j)$
- 4: NAssoc(C_k) \leftarrow NAssoc(C_{k+1}) + $\Delta(i^*, j^*)$
- 5: Update rows and columns of Δ , the weights, and the degrees corresponding to clusters i^* and j^* (3)–(5)
- 6: **end for**

MODEL ORDER SELECTION

- 1: **for** $k = n-1$ to 2 **do**
- 2: Calculate Curv(k) (6)
- 3: **end for**
- 4: $k^* \leftarrow \arg \max_k \text{Curv}(k)$ or provided by the user

REFINEMENT

- 1: Obtain flat partitioning for level k^*
- 2: Build the initial B and I (7), and refine \leftarrow **true**
- 3: **while** refine **do**
- 4: $\bar{\delta} \leftarrow 0$
- 5: **for** $u = 1$ to n **do**
- 6: **if** $u \in C_i$ is on the boundary and $\max_j \delta(u,i,j) > 0$ **then**
- 7: Move u from C_i to C_{j^*} , where $j^* = \arg \max_j \delta(u,i,j)$
- 8: Update B , I , the weights, and the degrees (9)–(18)
- 9: $\bar{\delta} \leftarrow \bar{\delta} + \delta(u,i,j^*)$
- 10: **end if**
- 11: **end for**
- 12: **if** $\bar{\delta} = 0$ **then**
- 13: refine \leftarrow **false**
- 14: **end if**
- 15: **end while**

4.4. Implementation and computational complexity of GANC

We take a similar approach to [57] when implementing the agglomerative clustering procedure of GANC. Max-heaps and balanced binary trees are used to store the rows of the Δ matrix and the adjacency matrix. A separate heap is also used to store the maximum of each row of Δ . This leads to the complexity of $O(mh \log(n))$ for the agglomerative clustering procedure, where h is the height of the generated dendrogram and m is the number of edges with non-zero weight (see [57] for details).

The model order selection step of the algorithm requires $O(n)$ computations. The computational requirements are much less than those of the methods analyzing the eigenvalue fall-off which require eigendecomposition (e.g., [49,47]). Performing an eigendecomposition to study the eigenvalue fall-off for all of the eigenvalues requires $O(n^3)$ operations.

To implement the refinement step, we use the map data structure to store and update B . Every row of B is stored in a map (C++ STL implementation of the map data structure is used [58]). Updates (9)–(12) are performed in constant time. To access each map entry of a row of B to insert, update, or delete, no more than $O(\log k)$ operations are required, because the maximum number of clusters connected to a boundary node is strictly less than k . Updates (13)–(15) each take no more than $O(\log k)$ operations and are repeated for all neighbors of a node that is moved. We have observed through experiments that nodes with larger degrees do not tend to be on the boundaries and hence the average number of neighbors to be updated is practically smaller than the average node degree. Eqs. (16) and (18), and any insertion to or deletion from the set of boundary nodes is performed in $O(\log k)$ time. Eq. (17) is trivial and is performed in a constant time. Not more than n nodes can be moved in a given iteration. When node i is moved, $O(d_i \log k)$ operations are required, where d_i is the degree of node i . Summing over all boundary nodes, the complexity is $O(m \log k)$ because $\sum_{i=1}^n d_i = 2m$. Hence the number of operations in a single iteration is at most $O(m \log k)$. Throughout our experiments we have observed very few refinement iterations even when the refinement is applied to the networks with millions of nodes and edges. Furthermore the algorithm is capable of limiting the number of iterations. Hence the complexity of the average number of operations of the whole refinement step does not exceed $O(m \log k)$.

The total computational complexity of GANC is dominated by the agglomerative clustering procedure. In practice, the height of the resulting dendrogram is often $O(\log n)$. Also, the graphs that are studied are usually sparse and hence $m = O(n)$. Therefore the complexity of the agglomerative clustering step of our algorithm is $O(n \log^2 n)$ for many real-life graphs.

The worst case runtime happens when the graph is complete and the dendrogram is totally unbalanced. There are $n-1$ agglomerations to construct the whole hierarchy each of which requires at most $O(n \log n)$ operations. Hence the worst case complexity for the agglomerative clustering step of the algorithm is $O(n^2 \log n)$, but this not expected to occur in practical situations.

The memory requirement of GANC is $O(m)$ which is equivalent to $O(n)$ for sparse graphs. In summary, in many practical cases (relatively sparse graphs and balanced dendrograms), the computational complexity is $O(n \log^2 n)$ and the memory requirement is $O(n)$. GANC can be downloaded from <http://www.ece.mcgill.ca/~coates/software/>.

5. Experimental results

In this section we provide a comparison of the performance and the runtime of our proposed algorithm GANC with a selection of the state-of-the-art graph clustering algorithms from the

literature. Implementations were downloaded from the corresponding authors' websites. Experiments were performed on an Intel 3.0 GHz Core 2 Quad CPU with 8 GB of RAM and Ubuntu 9.10 operating system.

5.1. Comparing algorithms

We compare to the following algorithms: Shi and Malik (recursive NCut) [9]; Meila and Shi (k-way NCut) [18]; Ng et al. [19]; Dhillon et al. [11]; Rosvall and Bergstrom [7]; and Blondel et al. [56]. The first four algorithms focus on maximizing NAssoc. Despite not addressing our criterion of interest directly, the other algorithms are included because they are scalable and represent the state-of-the-art in graph clustering.

The discussed clustering algorithms [9,18,19] are not scalable as they include eigendecomposition. Dhillon, Guan, and Kulis proposed an algorithm that strives to maximize normalized association without requiring any eigendecomposition [11]. Similar to the other algorithms that address the maximization of normalized association, the Dhillon–Guan–Kulis algorithm requires the user to provide the number of clusters. After specifying the number of clusters, three steps are performed: a coarsening phase, a base clustering step using region growing, and finally a refinement step using weighted kernel k -means with appropriate choices of the kernel and the node weights to maximize the normalized association.

Minimization of the criterion proposed by Rosvall and Bergstrom [7] results in a coarse-grained representation of the information flow through the network. Their proposed clustering objective aims to compress the underlying random walk on a graph without losing important network structures. To facilitate this, a coding is designed as follows: every cluster is assigned a unique code; in a given cluster, every node is also assigned a unique code. However two nodes in different clusters are allowed to be assigned to the same code. The algorithm then strives to minimize the average description length of a single step of a random walk.

The algorithm by Blondel et al. targets maximizing modularity [56]. Modularity (Q) is defined as the total fraction of intra-cluster sum-weight minus the expected fraction if the edges (and weights) were distributed randomly while the node degrees were preserved:

$$Q = \sum_{i=1}^k \left(\frac{w(C_i, C_i)}{M} - \frac{d(C_i)^2}{M^2} \right), \quad (19)$$

where $M = \sum_i \sum_j w(i, j)$. Modularity provides a valuable metric of the connectedness of clusters, but a number of authors have demonstrated that it suffers from a resolution limit when used to select the number of clusters [59]. The partitioning that maximizes modularity will generally not isolate clusters if the number of edges in the cluster is a small fraction of the total number of edges in the graph. The reason behind the resolution limit of modularity is the second term in the summation of (19); when the network gets larger, M increases monotonically. However the cluster degrees are not necessarily a function of the network size and are often bounded, regardless of the number of nodes [60]. This results in $d(C_i)^2/M^2 \ll 1$ and hence the value of modularity becomes dependent only on $w(C_i)/M$. By normalizing the summation by $d(C_i)$ as suggested in [12], the resolution limit phenomena is resolved; but we have

$$Q_{\text{normalized}} = \sum_{i=1}^k \frac{1}{d(C_i)} \left(\frac{w(C_i, C_i)}{M} - \frac{d(C_i)^2}{M^2} \right) = \frac{1}{M} [\text{NAssoc}(C_k) - 1], \quad (20)$$

i.e., maximization of $Q_{\text{normalized}}$ and normalized association are equivalent.

5.2. Synthetic graphs

We first analyze the performance on synthetic graphs, for which the true clustering behavior is known. We use benchmark graphs developed by Lancichinetti, Fortunato, and Radicchi [61] (LFR graphs). These random graphs are designed based on the planted partition model [62]. Each node is assigned to one of the k clusters. When edges are added to the graph, the probability that the edge is between nodes in the same cluster is $1 - \mu$ and the probability that the edge joins nodes from different clusters is μ . The LFR benchmarks have heterogeneous cluster sizes with user-specified lower bound and upper bound, c_{\min} and c_{\max} , respectively. Furthermore the node degrees are upper-bounded by d_{\max} and the average node degree is denoted by d_{avg} . As μ decreases, edges are increasingly likely to be intra-cluster, making the partitioning task easier.

The Meila–Shi algorithm performs at least as well as the other spectral clustering algorithms (Ng–Jordan–Weiss [19] and Shi–Malik [9]) and hence we only display the Meila–Shi results. The local search of Dhillon–Guan–Kulis is also excluded often, because it results in negligible improvement while it introduces a very large computational overhead.

5.2.1. Maximizing normalized association

Fig. 1 examines how the algorithms perform with respect to maximizing NAssoc for 1000-node LFR graphs. In order to observe how the algorithms perform on graphs with heterogeneous clusters, we let the cluster sizes range from 20 to 50 nodes.⁵ For each value of μ , 100 graph realizations are generated. The value of NAssoc is divided by k to obtain a value between 0 and 1 which represents the average NAssoc per cluster.

The algorithms perform almost identically with the exception of the Dhillon et al. algorithm [11]. The refinement step of GANC results in a significant improvement in the value of NAssoc. However the local search of Dhillon's algorithm does not improve NAssoc significantly. Note that the Blondel et al. algorithm results in higher values of NAssoc due to selecting smaller number of clusters than the true ones ($\text{NAssoc}(C_k)/k$ is a monotonically decreasing function of k); a fair comparison between algorithms is possible only when the number of clusters are equal.

5.2.2. Comparing to the planted partitions

The advantage of exploring performance on synthetic graphs is that a ground truth partitioning is available. Because the LFR benchmarks are based on the planted partition model, there is the additional advantage that NAssoc is an appropriate criterion to adopt. Apart from some possible small errors due to the randomness inherent in the construction of the benchmark graphs, the ground truth partitioning will correspond to a maximum of NAssoc for a given value of k .

For comparing two partitions on the same graph, we use the Jaccard index [64] which for two partitions, \mathcal{X} and \mathcal{Y} , is defined as

$$Jl = \frac{a}{(a+b+c)}, \quad (21)$$

where a , b , and c are, respectively, the number of pairs of nodes that are assigned to the same cluster in both \mathcal{X} and \mathcal{Y} , the same cluster in \mathcal{X} and different clusters in \mathcal{Y} , and the same cluster in \mathcal{Y} and different clusters in \mathcal{X} . When two partitions are identical, the Jaccard index evaluates to one, and it decreases as the partitions deviate from each other.

⁵ In the real-life networks that we consider in this paper, the clusters have been observed to be of a limited size, regardless of the network size [60]; for example the Dunbar number suggests an upper limit of 150 nodes for clusters in a social network [63].

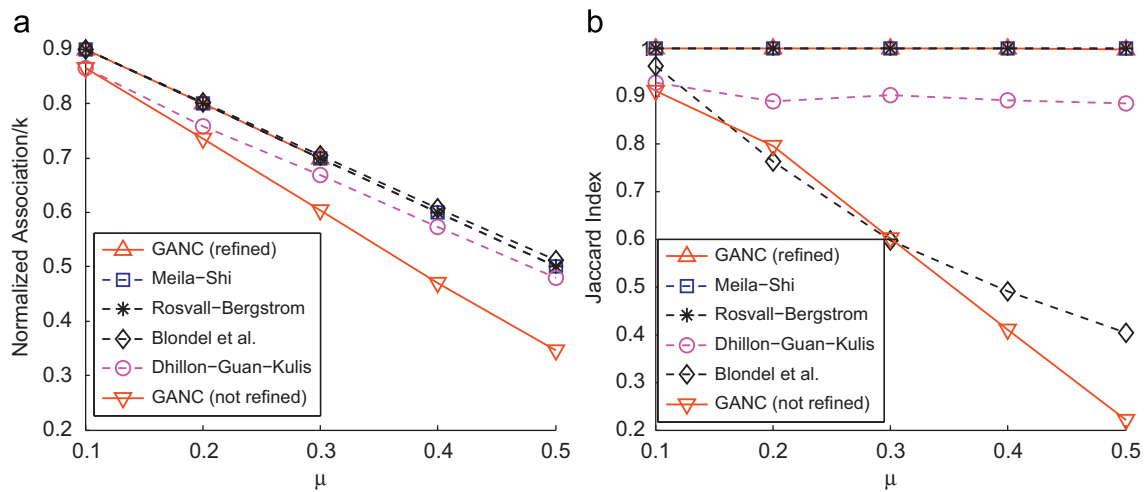


Fig. 1. LFR benchmarks with $c_{min} = 20$, $c_{max} = 50$, $d_{max} = 30$, $d_{avg} = 25$, and variable proportions of inter-cluster edges (μ). In all of the above figures, Meila-Shi, Rosvall-Bergstrom, and GANC (refined) are overlapping. In (b) and (d) Jaccard index=1 for the overlapping algorithms that corresponds to perfect extraction of the clusters: (a) $n=5000$ and (b) $n=5000$.

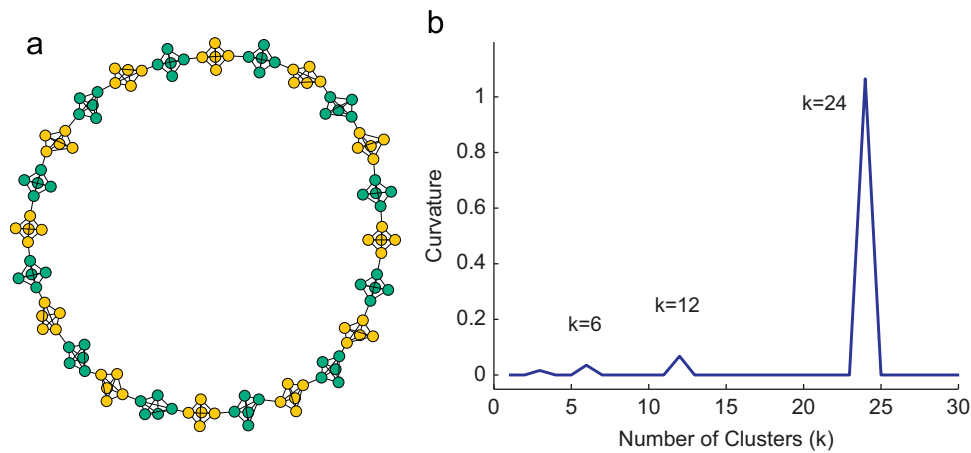


Fig. 2. The ring of 24 cliques; the highest peak of curvature is at 24: (a) ring graph and (b) curvature plot of the ring.

Fig. 1b shows how closely the partitionings identified by the algorithms match the ground truth in terms of the Jaccard index. The algorithm by Blondel et al. deviates dramatically from the true clustering (Fig. 1b). The reason is that the minimum and maximum cluster sizes are fixed for all the networks in addition to the average node degree. Hence the average number of edges per cluster remains the same making the proportion of the intra-cluster edges decrease. Due to the resolution limit of modularity, as the proportion of intra-cluster edges is decreased, modularity maximization algorithms tend to group several clusters into a single cluster. If we reduce n to 1000, the Blondel et al. algorithm leads to closer results to the ground truth. The algorithm by Dhillon et al. does not perform as well as other algorithms that strive to maximize NAssoc and the local search results in negligible improvement in terms of the Jaccard index.

5.3. Model order selection: the curvature metric

We first provide an illustrative example of the use of the curvature metric for selecting the number of clusters in the partitioning. This example highlights the difference in behavior

compared to the modularity metric used by modularity maximization algorithms [56].

An example used by Good et al. [59], to illustrate the resolution limit is the ring of 24 cliques,⁶ each of which has five nodes and is connected to its neighboring clique by a single edge. The value of modularity is maximized by a 12-cluster partition that merges pairs of cliques together. A more natural clustering is to consider each clique as an individual cluster. Fig. 2b shows the curvature plot for the case of the ring of cliques. In addition to the peak of the curvature that is correctly located at 24 clusters, it is interesting to note that the other local peaks of the curvature are also meaningful; e.g., the peak at 12 clusters corresponds to the clustering that maximizes modularity.

To see how the curvature indicates the true number of clusters for networks with more complex structures, we consider the LFR benchmarks. We fix the LFR parameters for a 1000-node graph and explore how the curvature changes as μ is varied. Fig. 3 indicates that the peak of the curvature plot correctly identifies the true number of clusters for μ in the range 0.1–0.5. As μ

⁶ A clique is a subset of nodes in a graph which are fully connected.

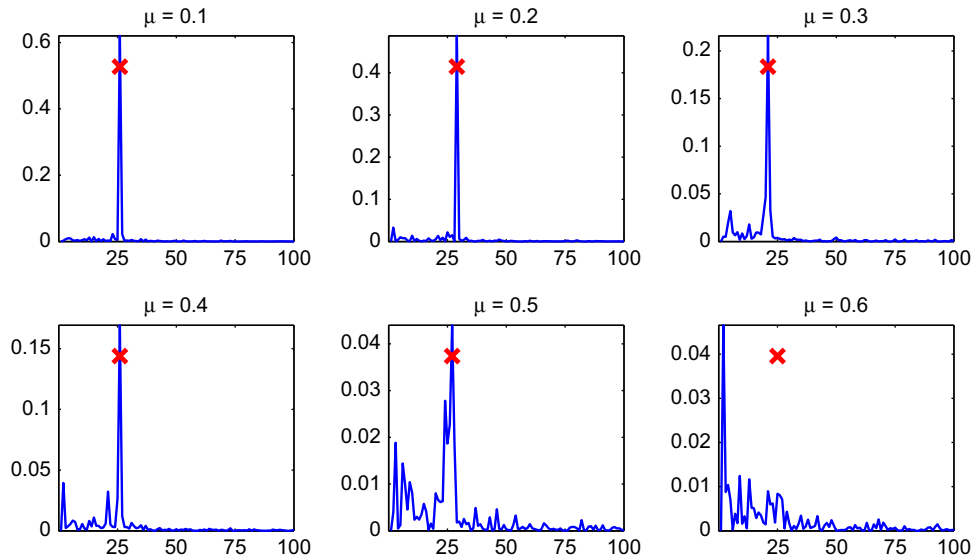


Fig. 3. Curvature versus number of clusters for varying proportion of inter-cluster edges (μ). The true clustering indicated by red \times . The peak of the curvature becomes less and less distinguishable by increasing μ . LFR benchmark graph with 1000 nodes and $c_{min} = 20, c_{max} = 50, d_{max} = 30, d_{avg} = 25$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

increases, the main peak becomes less distinct; for $\mu = 0.5$, the second highest peak is almost as large as the primary peak. As the clusters become increasingly interconnected, the curvature plot provides a less clear indication of the true number of clusters.

Note that Fig. 3 corresponds to a single realization of the LFR benchmarks. The empirical probabilities that the highest peak of the curvature corresponds to the true number of clusters for 100 realizations of the LFR benchmark with parameters as in Fig. 3 are $p_1 = [1.0, 1.0, 0.96, 0.60, 0.42, 0.06]$, for $\mu = [0.1, 0.2, 0.3, 0.4, 0.5, 0.6]$ when no refinement step is applied; furthermore the empirical probabilities that one of the two highest peaks corresponds to the true number of clusters are $p_2 = [1.0, 1.0, 1.0, 0.98, 0.75, 0.13]$.

These probabilities are obtained using the approximate curvatures derived from the agglomerative clustering procedure. To achieve a better insight into the model selection capabilities of the curvature metric, we derive better approximations by applying the refinement step to every level of the hierarchy and then recomputing curvature estimates. With this new procedure, we obtain $p_1 = [1.0, 0.99, 0.99, 0.95, 0.87, 0.77]$ and $p_2 = [1.0, 0.99, 0.99, 0.96, 0.88, 0.78]$. This indicates that curvature can provide a good indication of model order even for $\mu = 0.6$.

The model order selection algorithm of Rosvall–Bergstrom algorithm quite successfully indicates the true number of clusters in the case of LFR benchmarks; the empirical probabilities for the Rosvall–Bergstrom algorithm is $p_1 = [1.0, 1.0, 1.0, 0.98, 0.85, 0.83]$. The Blondel et al. algorithm does not successfully identify the true number of clusters when μ increases: $p_1 = [1.0, 0.99, 0.98, 0.79, 0.07, 0.0]$. Note that when $n=5000$, the failure of the Blondel et al. algorithm becomes more evident and p_1 turns to a 0% success for any μ (this is reflected in Fig. 1). However the refined GANC and Rosvall–Bergstrom algorithms obtain the same success ratios when $n=5000$.

5.4. Real networks

In this section, we examine the behavior the GANC and compare it with the other algorithms for graphs representing real networks. We first examine the behavior for small networks where there is knowledge of the ground truth partition. We then experiment with large networks, which allows us to assess the scalability of GANC.

Table 1

Small real networks: k known in advance (average NAssoc per cluster/Jaccard index). Maximum NAssoc is indicated by bold text.

Algorithm	Karate ($k=2$)	Football ($k=11$)	PolBooks ($k=3$)
GANC (refined)	0.872 /0.89	0.704/0.820	0.88 /0.67
Dhillon–Guan–Kulis	0.820/0.64	0.700/0.820	0.87/0.64
Ng–Jordan–Weiss	0.872 /1.00	0.701/0.825	0.82/0.57
Meila–Shi	0.869/0.89	0.696/0.820	0.88 /0.67
Shi–Malik	0.872 /1.00	0.706 /0.825	0.86/0.67

Table 2

Small real networks: Jaccard index for the estimated model orders.

Algorithm	Karate	Football	PolBooks
GANC (curv.)	0.80 ($k=3$)	0.76 ($k=13$)	0.69 ($k=2$)
Blondel et al.	0.53 ($k=4$)	0.70 ($k=11$)	0.67 ($k=3$)
Rosvall–Bergstrom	0.71 ($k=3$)	0.83 ($k=12$)	0.63 ($k=5$)

5.4.1. Small networks

We conduct experiments using the Zachary karate club network [65], the football network [6], and the political books network.⁷ The Zachary karate club network portrays friendships in a karate club. During the polling period, there was a dispute between the manager and the instructor which led to the establishment of a new club by the instructor. The karate students then split into two groups, either staying with the original club or following the instructor to the new club. The two clusters associated with the network correspond to these two groups. Each node in the football networks corresponds to a team in US college football. There are 11 conferences and five independent teams. Each edge corresponds to the existence of a match between the connected nodes. The independent teams can be interpreted as outliers. Nodes of the political book network correspond to the political books sold by Amazon.com and are categorized as neutral, conservative, and liberal. The edges between pairs of nodes correspond to frequent co-buying of the pairs of books by the same buyers. The three networks are unweighted.

⁷ Collected by Krebs, <http://www.orgnet.com>.

Table 3

NAssoc of the cortical networks. Maximum NAssoc is indicated by bold text. The model order selected for each column in each subtable is identified within parentheses.

Algorithm	A-1	A-2	B	C	D	E
Rosvall–Bergstrom	0.725 (62)	0.706 (68)	0.738 (61)	0.706 (62)	0.691 (63)	0.706 (58)
GANC (refined)	0.730	0.716	0.740	0.712	0.705	0.718
Dhillon–Guan–Kulis	0.696	0.682	0.712	0.688	0.670	0.684
Shi–Malik	0.715	0.692	0.727	0.690	0.686	0.694
Blondel et al.	0.873 (14)	0.872 (13)	0.885 (15)	0.862 (14)	0.856 (14)	0.828 (21)
GANC (refined)	0.886	0.879	0.883	0.873	0.864	0.840
Dhillon–Guan–Kulis	0.860	0.865	0.872	0.859	0.851	0.810
Shi–Malik	0.877	0.874	0.877	0.860	0.870	0.826
GANC (refined)	0.908 (10)	0.945 (4)	0.899 (12)	0.880 (13)	0.930 (5)	0.907 (9)
Dhillon–Guan–Kulis	0.889	0.946	0.897	0.863	0.916	0.885
Shi–Malik	0.898	0.951	0.886	0.863	0.936	0.902

Tables 1 and 2 compare the performance of GANC with other clustering algorithms for these small, well-known datasets. The tables provide the average NAssoc per cluster (NAssoc/ k) and the Jaccard index. The “true” number of clusters is indicated within parentheses beside the name of each network.⁸ In Table 1, all algorithms (including GANC) are informed of the true number of clusters and instructed to generate a partitioning with that number of clusters. Table 1 indicates that the spectral clustering algorithms (Ng–Jordan–Weiss, Meila–Shi, and Shi–Malik) achieve similar performance in terms of NAssoc. We note that when k is specified for these three networks, GANC identifies partitionings with average NAssoc as large or larger than those of the partitionings identified by the spectral clustering techniques, with the exception of Shi–Malik in the case of the football network.

Table 2 compares the performance of algorithms that select the number of clusters based on some aspect of the data. The number of clusters selected by each algorithm is shown within parentheses after the Jaccard index. Rosvall’s algorithm uses a description length metric to select the number of clusters [7]; Blondel’s algorithm chooses the partitioning that maximizes the modularity [56]. A direct comparison of NAssoc is not valid when k is not fixed.

The first row of this table shows the performance of GANC when the curvature plot is used to select the number of clusters. Although the selected number of clusters does not correspond to the true number of clusters for any of the networks, all three values can be explained. In the karate club network there is a group of students who have weak ties to other members of the network and these are identified as a third cluster; in the football network, GANC isolates the independent teams; in the political books network, GANC does not identify a neutral cluster and assigns each of the neutral books to either the liberal or conservative cluster.

5.4.2. Cortical networks

Here we study the networks presented in [66] which are weighted graphs, each with 998 nodes. The nodes correspond to small regions of the human cerebral cortex and the edges correspond to cortico-cortical axonal pathways. The networks are developed for five patients (the extraction is performed twice for patient A). We first perform a comparison of the competing algorithms, then we discuss the clustering results of GANC.

When applying the Rosvall–Bergstrom and Blondel et al. algorithms the user does not have the freedom to choose the number of clusters. Hence we repeat the experiment twice to make a meaningful objective comparison in terms of NAssoc. Table 3 lists

the average NAssoc of every clustering algorithm and patient. In the first subtable, k is set to the model order selected by the Rosvall–Bergstrom algorithm. Then, in the second subtable, k is set to the model order selected by the Blondel et al. algorithm, and in the third subtable, peaks of the model order selected by curvature are used. Note that the Shi–Malik algorithm outperforms the other two spectral algorithms for these datasets; hence we only include the Shi–Malik results.

The clustering results corresponding to the curvature peaks include cluster(s) that contain nodes from both of the brain hemispheres and clusters that include nodes from only one of the hemispheres. In the following discussion, we call the clusters that contain nodes from both of the hemispheres the *central clusters*. Graphs A1, B, D, and E include only one central cluster. Graphs A2 and C each include two central clusters.

The regions of the brain that are commonly grouped in the central clusters are posterior cingulate cortex, precuneus, cuneus, paracentral lobule, pericalcarine cortex, caudal anterior cingulate cortex, isthmuscingulate, isthmus of the cingulate cortex, and lingual gyrus, except in the case of graph B. If the second largest peak of the curvature is selected for B ($k=5$), the same regions are assigned to its central cluster. The first five of the mentioned regions are also classified as part of the structural core proposed by Hagmann et al. [66]; Hagmann et al. used cluster strengths, cluster degrees, k -Core, s -Core, betweenness centrality, and efficiency⁹ to propose a structural core of the brain consisting of eight regions.

The authors of [66] also used modularity maximization and found six modules, two of which included nodes from both of the hemispheres (central clusters). The regions that are assigned to the central clusters by modularity maximization are similar to the ones extracted by GANC.

5.4.3. Larger networks

Here we illustrate the performance of our algorithm on large graphs. We apply our algorithm to four networks with different natures: Cond-Mat (a collaboration network) [67], Googleweb (a web graph) [60],¹⁰ Amazon (a product co-purchasing network) [68], and AS-Skitter (an autonomous system graph) [69]. The nodes in Cond-Mat represent scientists that submit their pre-prints to the condensed matter archive at www.arxiv.org; the

⁸ Note that the “true” number of clusters is a somewhat artificial social construct and does not necessarily correspond to a partitioning that maximizes any meaningful graph clustering metric.

⁹ Cluster strength and degree are the weighted and unweighted degrees, respectively. k -Core (s -Core) is the largest subgraph that includes nodes of degree (respectively, strength) at least k (respectively, s). Betweenness centrality of a region is a measure of the proportion of the shortest paths passing through it. Efficiency of a region indicates how short the region’s average path lengths to other regions are.

¹⁰ Google programming contest: <http://www.google.com/programming-contest/>.

Table 4
Larger real networks: Partition metrics ((Number of clusters) NAssoc per cluster/execution time).

Network	Cond-Mat	Googleweb	Amazon	AS-Skitter
# of Nodes/Edges	36458/171736	342408/1142134	403364/2443311	1694616/11094209
Avg/Max Degree	28.5/278	6.7/1367	12.1/2752	13.1/35455
GANC	(2121) 0.784/0m2s	(12213) 0.869/2m30s	(11237) 0.768/3m4s	(26399) 0.769/46m33s
Dhillon-Guan-Kulis	(2121) 0.669/0m3s	(12213) –	(11237) 0.679/6m0s	(26399) –
Rosvall-Bergstrom	(2121) 0.746/0m10s	(12213) 0.844/ 1m17s	(11237) 0.712/11m26s	(26399) 0.720/99m37s
GANC	(82) 0.957/0m2s	(239) 0.993/2m30s	(230) 0.984/3m4s	(1776) 0.933/46m33s
Dhillon-Guan-Kulis	(82) 0.784/ < 1s	(239) 0.935/ 0m2s	(230) 0.872/ 0m4s	(1776) 0.709/8m0s
Blondel et al.	(82) 0.849/0m1s	(239) 0.990/0m5s	(230) 0.952/0m14s	(1776) 0.859/ 0m4s
GANC	(37) 0.970/0m2s	(7) 0.999/2m30s	(22) 0.997/3m4s	(32) 0.994/46m33s
GANC	(526) 0.902/0m2s	(855) 0.984/2m30s	(206) 0.996/3m4s	–

edges represent co-authorships. The nodes in Googleweb represent websites and the directed edges represent the existence of hyperlinks. Each node in Amazon network corresponds to a product purchased at www.amazon.com. Each directed edge from a node to another means that when the former is purchased, the latter is frequently also purchased. AS-Skitter is a network of autonomous systems extracted by traceroute analysis.¹¹

The Cond-Mat graph is weighted and the rest of the graphs are unweighted. In Table 4, the average and maximum degrees do not take the weights into account; the number of edges affects the runtime of GANC, not the weight values. None of the above graphs are originally connected. However each has a very large connected component that includes the majority of the nodes and the edges. Here we conduct clustering analysis of the largest connected component. Some of these graphs are directed, but we construct undirected graphs by adding the adjacency matrix to its transpose.

Table 4 compares the performance of the algorithms that are scalable to such large networks. The table shows the average NAssoc per cluster of the identified partitioning and the time required for completion of the algorithm. The spectral clustering algorithms cannot be executed on our test machine when either the number of nodes or the number of clusters is very large. Both the computational time and the memory requirements are excessive. We therefore compare the performance of GANC, Dhillon-Guan-Kulis (no local search), Rosvall-Bergstrom, and the Blondel et al. algorithms.

The Rosvall-Bergstrom and Blondel et al. algorithms automatically select the number of clusters. A meaningful comparison of the average NAssoc is only possible when the number of clusters are equal; to facilitate comparisons, we therefore construct multiple partitionings for GANC and the Dhillon-Guan-Kulis algorithms, each with a different number of clusters. We also construct a partitioning corresponding to one or more of the peaks in the curvature plot (for some of the networks, there are two peaks that are very similar in value, so we consider it useful to examine the partitionings corresponding to each). The last two rows of Table 4 correspond to those peaks.

GANC is superior to other algorithms with respect to maximizing NAssoc in all of the cases. GANC also significantly outperforms the Dhillon-Guan-Kulis algorithm, which is also striving to maximize NAssoc. The latter is also outperformed by Rosvall-Bergstrom and Blondel et al. algorithms.

In terms of the computation time, GANC is often faster than Rosvall-Bergstrom, but the scaling behavior is different. To illustrate this, in addition to the graphs listed in Table 1, we have applied the algorithms to the road network graph of California [60] which is extremely sparse (the average degree is 2.8). The graph contains around two million nodes. While GANC performs

the clustering in 32 s, it takes 197 min for the Rosvall-Bergstrom algorithm to converge to a solution. The complexity of GANC is dominated by the agglomerative clustering procedure which requires $O(mh \log n)$ operations. Hence the speed of GANC depends on the number of edges in the graph, and the height of the dendrogram. However the graph sparsity does not show an impact on Rosvall-Bergstrom's execution time. Blondel's algorithm is much faster than GANC, but as mentioned previously, it suffers from the resolution limit associated with clustering algorithms that maximize modularity. The Dhillon-Guan-Kulis algorithm is also very fast for small values of k , but it gets slower and its memory requirements become excessive if k becomes large.

To illustrate the properties of the algorithm outputs in terms of the cluster sizes and the connectivity of individual clusters, we have examined the plots of NAssoc of individual clusters versus the number of nodes in them. Each algorithm behaves similarly on different graphs presented in Table 4. Fig. 4a compares GANC and Rosvall-Bergstrom when applying these algorithms on the Amazon co-purchasing graph. For the same number of clusters, GANC produces clusters with higher values of NAssoc than Rosvall-Bergstrom. The latter produces many small clusters with very low values of NAssoc. The behavior of the Dhillon-Guan-Kulis and the Blondel et al. algorithms are discussed in our case study.

5.5. Case study: US patent citation graph

As a case study, we consider the undirected version of the citation graph released by the National Bureau of Economic Research [70,71]. The patents are classified into six broad technological categories. A more refined classification leads to 36 sub-categories. We use each patent's label (category or sub-category) in addition to NAssoc and runtime to perform a comparison.

The original graph is not connected; however the largest connected component contains more than 99.7% of the nodes and all of the edges. Hence we focus on the largest connected component which contains 3,764,117 nodes and 16,511,740 edges. The maximum node degree is 793.

5.5.1. Clustering runtime

The Rosvall-Bergstrom algorithm [7] was terminated after more than 30 h without converging to a solution. GANC takes 77 min to construct the full hierarchy. After the hierarchy has been generated, each flat partitioning including the refinement takes less than 35 s. The algorithm by Blondel et al. [56] takes 8 min. The Dhillon-Guan-Kulis algorithm [11] takes 72 s for $k=57$ and its execution time increases as k is increased.

5.5.2. Maximization of normalized association

In order to have a fair comparison in terms of NAssoc, we fix $k=57$ to match the number selected by the Blondel et al.

¹¹ <http://www.caida.org/tools/measurement/skitter>

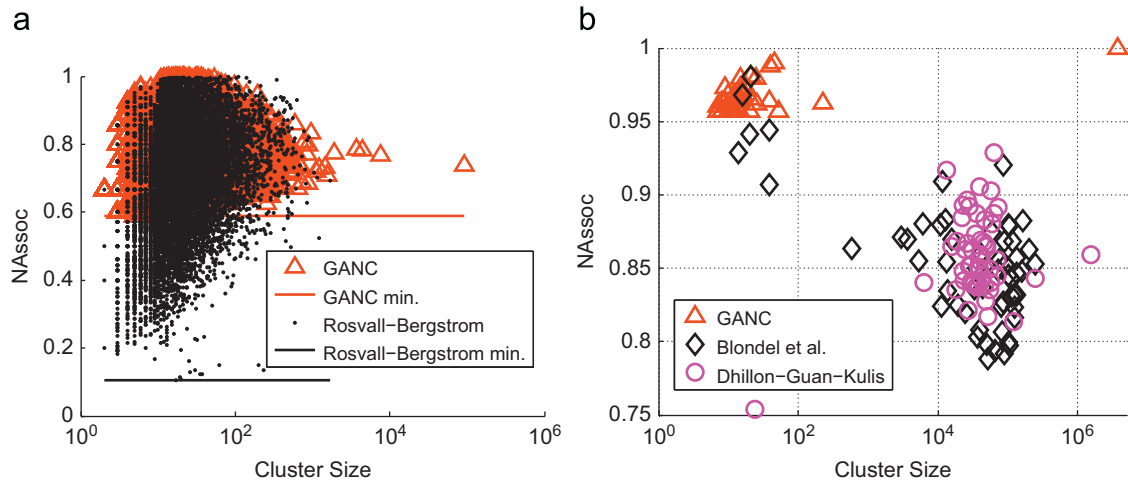


Fig. 4. NAssoc of individual clusters: (a) Amazon and (b) US patent citation.

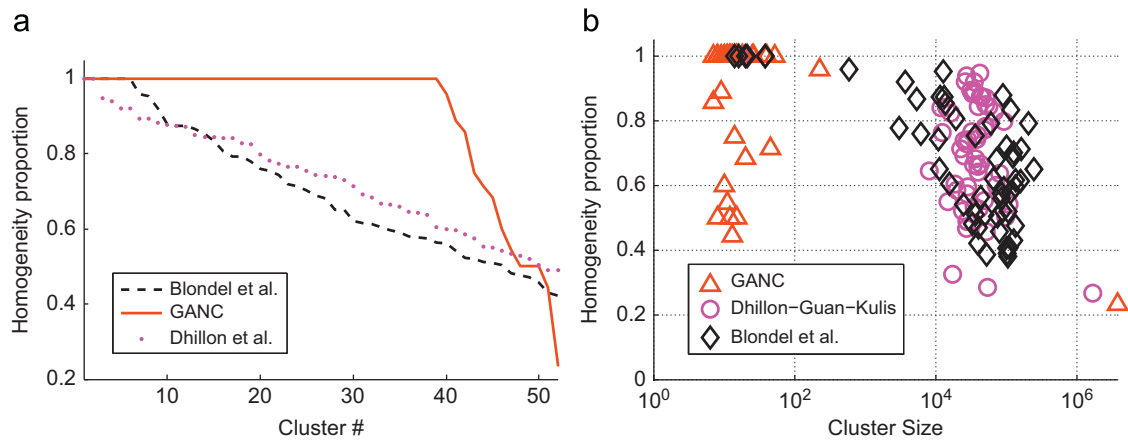


Fig. 5. The 52 most homogeneous clusters out of 57 clusters using the Blondel et al. and the Dhillon–Guan–Kulis algorithms: (a) maximum homogeneity proportions and (b) homogeneity proportion per cluster.

algorithm. The values of $NAssoc/k$ for GANC, Dhillon–Guan–Kulis, and Blondel et al. are 0.964, 0.859, and 0.855, respectively. The values of NAssoc of the individual clusters are shown in Fig. 4b which illustrates the clear superiority of GANC.

5.5.3. Extraction of true clusters and absence of large well-defined clusters

We use the categories and sub-categories to classify the nodes of the patent citation graph. We denote the maximum proportion of nodes from the same class in a given cluster as the *homogeneity proportion* of that cluster. For the Blondel et al. and Dhillon–Guan–Kulis algorithms, we use $k=57$ and for GANC we use $k=52$ (the closest peak of the curvature plot). The clusters are sorted according to their homogeneity proportion in Fig. 5a. The figure shows the superiority of GANC in extracting clusters of nodes from the same categories. When sub-categories are employed for the assessment, the superiority of GANC becomes more pronounced.

Fig. 5a alone does not provide a fair comparison as any singleton would have homogeneity proportion of 1. In Fig. 5b, the homogeneity proportion is plotted versus the size of the extracted clusters for the three algorithms.

The Blondel et al. algorithm produces clusters as small as 14 to as large as 250,000 nodes. However the largest homogeneous cluster that is extracted has 38 nodes. The cluster quality degrades as the

cluster size increases. This trend is likely due to the resolution limit of modularity.

Both GANC and Dhillon–Guan–Kulis produce a very large cluster (the core cluster [60]). Excluding the core, the Dhillon–Guan–Kulis algorithm produces clusters of an average size of n/k which is around 37,000 nodes in this example (see Fig. 5b). This behavior of the Dhillon–Guan–Kulis algorithm is due to the underlying region growing procedure it adopts from Metis [17].

The clusters extracted by GANC are much smaller than the other two algorithms. Excluding the core, the cluster sizes are between 7 and 225. Similar to previous observations of Leskovec et al. [60], the strongly connected clusters of the large graphs listed in Table 4 diminish in number as we move towards the center of the graph. Having a small homogeneity proportion is then expected for the core as it includes several clusters that are not very well-connected and are lumped together. If we study a clustering that is in a lower level in the hierarchy, further clusters would be extracted from the core and the core shrinks. The resulting clusters are not as well-connected in terms of NAssoc though. In other words, the most isolated and well-connected clusters are extracted at the higher levels of the hierarchy.

6. Conclusion

We proposed a novel algorithm to maximize normalized association that consists of three steps: the agglomerative hierarchical

clustering procedure, the model order selection step, and the refinement step.

The agglomerative clustering procedure requires $O(n \log^2 n)$ operations for many real-life graphs, where n is the number of nodes in the graph. This procedure dominates the computational complexity of the other steps of the algorithm. The second step of our algorithm, the model order selection, is based on the relative improvement of normalized association when changing the number of clusters; the *curvature* plot is used to select one or several model orders for the final clustering. Unlike modularity, the curvature metric does not exhibit an intrinsic resolution limit. For a multi-resolution analysis, a user can specify a range on the allowable number of clusters, and the algorithm will select the number of clusters with the maximum curvature in that range. After selecting the number of clusters, the clusters are passed to the refinement step. This step of our algorithm iterates over the boundary nodes in the clusters and explores possible improvements in normalized association by moving each of the boundary nodes to their neighboring clusters. Using the map data structure, the overhead added by the refinement becomes negligible. Experiments show that despite the negligible runtime of the refinement step, it significantly improves the initial results with respect to the normalized association maximization.

Our experimental analysis on relatively small networks indicated that the proposed algorithm identified partitions that have normalized association values comparable to the spectral algorithms that involve an eigendecomposition. These algorithms are too computationally complex to be applied to very large graphs. We demonstrated that our proposed algorithm can be applied to large graphs (millions of nodes and edges). For these large graphs, the proposed algorithm identifies partitions that have larger values of normalized association than those identified by the only comparable algorithm that directly addresses the normalized cut metric.

A clustering algorithm can by no means be suitable for every application. For example, despite the failure of Dhillon–Guan–Kulis algorithm to extract clusters with nodes of the same categories in the patent citation network, it generates clusters of very similar sizes. This is more suitable for VLSI applications for instance [72]. On the other hand, when clustering is meant to extract “communities” (group of nodes with strong intra-connection), GANC and Rosvall–Bergstrom are clearly preferred.

The Blondel et al. and Rosvall–Bergstrom algorithms are not able to generate a partitioning to an arbitrary number of clusters¹²; even though they are agglomerative, they do not generate the complete hierarchy because they merge several nodes/clusters in each of their iterations. Hence if one is interested in several arbitrary clustering levels, GANC fits one’s requirement the best out of the competing algorithms discussed.

In this paper we only considered undirected graphs while the edge directions could carry valuable information about the structure of a graph. An extension of GANC can be developed by adopting the *generalized normalized cut* criterion [73].

References

- [1] S. Wasserman, K. Faust, *Social Network Analysis: Methods and Applications*, Cambridge University Press, 1994.
- [2] R. Guimera, L. Amaral, Functional cartography of complex metabolic networks, *Nature* 433 (7028) (2005) 895–900.
- [3] Y. Loewenstein, E. Portugaly, M. Fromer, M. Linial, Efficient algorithms for accurate hierarchical clustering of huge datasets: tackling the entire protein space, *Bioinformatics* 24 (13) (2008) i41–i49.
- [4] Z. Chen, Y. He, P. Rosa-Neto, J. Germann, A. Evans, Revealing modular architecture of human brain structural networks by using cortical thickness from MRI, *Cerebral Cortex* 18 (10) (2008) 2374–2381.
- [5] W. Crum, Spectral Clustering and label fusion for 3D tissue classification: Sensitivity and consistency analysis, in: *Proceedings of Medical Image Understanding and Analysis*, Dundee, Scotland, 2008.
- [6] M. Girvan, M.E.J. Newman, Community structure in social and biological networks, *Proc. Nat. Acad. Sci. USA* 99 (12) (2002) 7821–7826.
- [7] M. Rosvall, C.T. Bergstrom, Maps of random walks on complex networks reveal community structure, *Proc. Nat. Acad. Sci. USA* 105 (4) (2008) 1118–1123.
- [8] S. Lafon, A. Lee, Diffusion maps and coarse-graining: a unified framework for dimensionality reduction, graph partitioning, and data set parameterization, *IEEE Trans. Pattern Anal. Mach. Intell.* 28 (9) (2006) 1393–1403.
- [9] J. Shi, J. Malik, Normalized cuts and image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (8) (2000) 888–905.
- [10] C. Fowlkes, S. Belongie, F. Chung, J. Malik, Spectral grouping using the Nyström method, *IEEE Trans. Pattern Anal. Mach. Intell.* 26 (2) (2004) 214–225.
- [11] I. Dhillon, Y. Guan, B. Kulis, Weighted graph cuts without eigenvectors: a multilevel approach, *IEEE Trans. Pattern Anal. Mach. Intell.* 29 (11) (2007) 1944–1957.
- [12] L. Yu, C. Ding, Network community discovery: solving modularity clustering via normalized cut, in: *Proceedings of ACM Workshop Mining and Learning with Graphs*, 2010.
- [13] A. Sinclair, M. Jerrum, Approximate counting, uniform generation and rapidly mixing Markov chains, *Inform. Comput.* 82 (1) (1989) 93–133.
- [14] W. Donath, A. Hoffman, Algorithms for partitioning of graphs and computer logic based on eigenvectors of connection matrices, *IBM Tech. Disc. Bull.* 15 (3) (1972) 938–944.
- [15] A. Pothén, H. Simon, K.-P. Liou, Partitioning sparse matrices with eigenvectors of graphs, *SIAM J. Matrix Anal. Appl.* 11 (1) (1990) 430–452.
- [16] B. Hendrickson, R. Leland, A multilevel algorithm for partitioning graphs, in: *Proceedings of ACM International Conference on Supercomputers*, Barcelona, Spain, 1995.
- [17] G. Karypis, V. Kumar, A fast and high quality multilevel scheme for partitioning irregular graphs, *SIAM J. Sci. Comput.* 20 (1) (1998) 359–392.
- [18] M. Meila, J. Shi, A random walks view of spectral segmentation, in: *Proceedings of International Workshop on Artificial Intelligence Statistics*, Key West, FL, USA, 2001.
- [19] A. Ng, M. Jordan, Y. Weiss, On spectral clustering: Analysis and an algorithm, in: *Proceedings of Advanced Neuro Information Processing System*, Vancouver, BC, Canada, 2001.
- [20] D. Yan, L. Huang, M. I. Jordan, Fast approximate spectral clustering, in: *Proceedings of ACM International Conference on Knowledge Discovery and Data Mining*, Paris, France, 2009.
- [21] E. Sharon, A. Brandt, R. Basri, Fast multiscale image segmentation, in: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Hilton Head, SC, USA, 2000.
- [22] C. Ding, X. He, H. Zha, M. Gu, H. Simon, A min-max cut algorithm for graph partitioning and data clustering, in: *Proceedings of the IEEE International Conference on Data Mining*, San Jose, CA, USA, 2001.
- [23] E. Sharon, A. Brandt, R. Basri, Segmentation and boundary detection using multiscale intensity measurements, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Kauai, HI, USA, 2001.
- [24] A. Akselrod-Ballin, M. Galun, R. Basri, A. Brandt, M. Gomori, M. Filippi, P. Valsasina, An integrated segmentation and classification approach applied to multiple sclerosis analysis, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, New York, NY, USA, 2006.
- [25] J. Corso, E. Sharon, S. Dube, S. El-Saden, U. Sinha, A. Yuille, Efficient multilevel brain tumor segmentation with integrated Bayesian model classification, *IEEE Trans. Med. Imag.* 27 (5) (2008) 629–640.
- [26] G. Milligan, M. Cooper, An examination of procedures for determining the number of clusters in a data set, *Psychometrika* 50 (2) (1985) 159–179.
- [27] M. Halkidi, Y. Batistakis, M. Vazirgiannis, On clustering validation techniques, *J. Intell. Inf. Syst.* 17 (2) (2001) 107–145.
- [28] T. Caliński, J. Harabasz, A dendrite method for cluster analysis, *Commun. Stat. Theory Methods* 3 (1) (1974) 1–27.
- [29] L. Hubert, J. Levin, A general statistical framework for assessing categorical clustering in free recall, *Psychol. Bull.* 83 (6) (1976) 1072–1080.
- [30] F. Baker, L. Hubert, Measuring the power of hierarchical cluster analysis, *J. Am. Stat. Assoc.* (1975) 31–38.
- [31] W. Sarle, *Cubic Clustering Criterion*, Technical Report, Cary, NC: SAS Inst. Inc, 1983.
- [32] R. Tibshirani, G. Walther, T. Hastie, Estimating the number of clusters in a data set via the gap statistic, *J. R. Stat. Soc.: Series B* 63 (2) (2001) 411–423.
- [33] J. Edachery, A. Sen, F. Brandenburg, Graph clustering using distance-k cliques, *Graph Drawing*, Lecture Notes in Computer Science, vol. 1731, Springer, Berlin/Heidelberg, 1999, pp. 98–106.
- [34] R. Gnanadesikan, J. Kettenring, J. Landwehr, Interpreting and assessing the results of cluster analyses, *Bull. Int. Stat. Inst.* 47 (2) (1977) 451–463.
- [35] W. Krzanowski, Y. Lai, A criterion for determining the number of groups in a data set using sum-of-squares clustering, *Biometrics* 44 (1) (1988) 23–34.
- [36] T. Pedersen, A. Kulkarni, Automatic cluster stopping with criterion functions and the gap statistic, in: *Proceedings of Conference on North American Chapter on Association for Computational Linguistics and Human Language Technology*, New York City, NY, USA, 2006.

¹² The available implementation of the Rosvall–Bergstrom algorithm is based on the Blondel et al. algorithm. See <http://www.tp.umu.se/~rosvall/algorithm.pdf>.

- [37] S. Arnold, A test for clusters, *J. Market. Res.* 16 (4) (1979) 545–551.
- [38] G. Schwarz, Estimating the dimension of a model, *Ann. Stat.* 6 (2) (1978) 461–464.
- [39] H. Akaike, A new look at the statistical model identification, *IEEE Trans. Auto. Control* 19 (6) (1974) 716–723.
- [40] J. Wolfe, Pattern clustering by multivariate mixture analysis, *Multivar. Behav. Res.* 5 (3) (1970) 329–350.
- [41] D. Pelleg, A. Moore, X-means: Extending K-means with Efficient Estimation of the Number of Clusters, in: *Proceedings of the International Conference on Machine Learning*, Stanford, CA, USA, 2000.
- [42] M. Mariadassou, S. Robin, C. Vacher, Uncovering latent structure in valued graphs: a variational approach, *Ann. Appl. Stat.* 4 (2) (2010) 715–742.
- [43] P. Latouche, E. Birmelé, C. Ambroise, Bayesian methods for graph clustering, Technical Report, Laboratoire Statistique et Génome, UMR CNRS 8071-INRA 1152-UEVE, 91000 Evry, France, 2008.
- [44] P. Latouche, E. Birmelé, C. Ambroise, Bayesian methods for graph clustering, in: *Advances in Data Analysis Data Handling and Business Intelligence*, 2010, pp. 229–239.
- [45] Y. Gdalyahu, D. Weinshall, M. Werman, Stochastic image segmentation by typical cuts, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Fort Collins, CO, USA, 1999.
- [46] D. Harel, Y. Koren, On clustering using random walks, in: *Proceedings of Conference on Foundations of Software Technology and Theoretical Computer Science*, Bangalore, India, 2001.
- [47] A. Azran, Z. Ghahramani, Spectral methods for automatic multiscale data clustering, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, New York, NY, USA, 2006.
- [48] U. von Luxburg, A tutorial on spectral clustering, *Stat. Comput.* 17 (4) (2007) 395–416.
- [49] L. Zelnik-Manor, P. Perona, Self-tuning spectral clustering, in: *Proceedings of the Advanced Neural Information Processing Systems*, Vancouver, BC, Canada, 2004.
- [50] A. Ben-Hur, A. Elisseeff, I. Guyon, A stability based method for discovering structure in clustered data, in: *Pacific Symposium on Biocomputing*, Lihue, HI, USA, 2002.
- [51] T. Lange, V. Roth, M. Braun, J. Buhmann, Stability-based validation of clustering solutions, *Neural Comput.* 16 (6) (2004) 1299–1323.
- [52] S. Ben-David, U. Von Luxburg, D. Pál, A sober look at clustering stability, *Learn. Theory* (2006) 5–19.
- [53] M. Smolkin, D. Ghosh, Cluster stability scores for microarray data in cancer studies, *BMC Bioinformatics* 4 (1) (2003) 36.
- [54] S. I. Daitch, J. A. Kelner, D. A. Spielman, Fitting a graph to vector data, in: *Proceedings of the ACM International Conference on Machine Learning*, Montreal, QC, Canada, 2009.
- [55] M. Maier, U. Von Luxburg, M. Hein, Influence of graph construction on graph-based clustering measures, in: *Proceedings of the Advanced Neural Information Processing Systems*, Vancouver, BC, Canada, 2009.
- [56] V. Blondel, J. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech.: Theor. Exp.* 2008 (2008) P10008.
- [57] A. Clauset, M. Newman, C. Moore, Finding community structure in very large networks, *Phys. Rev. E* 70 (6) (2004) 66111.
- [58] P. Plauger, M. Lee, D. Musser, A. Stepanov, C++ Standard Template Library, Prentice Hall, PTR Upper Saddle River, NJ, USA, 2000.
- [59] B.H. Good, Y.-A. de Montjoye, A. Clauset, Performance of modularity maximization in practical contexts, *Phys. Rev. E* 81 (4) (2010) 46106–46124.
- [60] J. Leskovec, K. J. Lang, A. Dasgupta, M. W. Mahoney, Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters, *Comp. Res. Repos. abs/0810.1355*.
- [61] A. Lancichinetti, S. Fortunato, Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities, *Phys. Rev. E* 80 (1) (2009) 16118.
- [62] A. Condon, R.M. Karp, Algorithms for graph partitioning on the planted partition model, *Rand. Struct. Alg.* 18 (2) (2001) 116–140.
- [63] R. Dunbar, *Grooming, Gossip, and the Evolution of Language*, Harvard University Press, MA, USA, 1998.
- [64] M. Downton, T. Brennan, Comparing classifications: an evaluation of several coefficients of partition agreement, in: *Proceedings Meeting of the Classification Society*, Boulder, CO, USA, 1980.
- [65] W. Zachary, An information flow model for conflict and fission in small groups, *J. Anthropol. Res.* 33 (4) (1977) 452–473.
- [66] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. Honey, V. Wedeen, O. Sporns, Mapping the structural core of human cerebral cortex, *PLoS Biol.* 6 (7) (2008) e159.
- [67] M. Newman, The structure of scientific collaboration networks, *Proc. Nat. Acad. Sci. USA* 98 (2) (2001) 404–409.
- [68] J. Leskovec, L. A. Adamic, B. A. Huberman, The dynamics of viral marketing, in: *Proceedings of the International Conference on Electronic Commerce*, Ann Arbor, MI, USA, 2006.
- [69] J. Leskovec, J. Kleinberg, C. Faloutsos, Graphs over time: densification laws, shrinking diameters and possible explanations, in: *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, Chicago, Illinois, USA, 2005.
- [70] The NBER Citations Data File <<http://www.nber.org/patents/>>.
- [71] B. Hall, A. Jaffe, M. Trajtenberg, The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools, NBER Working Paper 8498.
- [72] J. Li, L. Behjat, A connectivity based clustering algorithm with application to VLSI circuit partitioning, *IEEE Trans. Circ. Syst. II: Exp. Briefs* 53 (5) (2006) 384–388.
- [73] M. Meila, W. Pentney, Clustering by weighted cuts in directed graphs, in: *Proceedings of the SIAM International Conference on Data Mining*, Minneapolis, MN, USA, 2007.

Seyed Salim Tabatabaei received the B.Sc. degree in computer engineering from Ajman University of Science and Technology in 2008, and the M.Eng. degree in electrical engineering from McGill University in 2011.

Mark J. Coates received the B.E. degree from the University of Adelaide in 1995 and the Ph.D. degree from the University of Cambridge in 1999. He was a research associate and lecturer at Rice University from 1999 to 2001. He joined McGill University in 2002 and is an Associate Professor.

Michael G. Rabbat earned the B.Sc. from the University of Illinois at Urbana-Champaign (2001), the M.Sc. degree from Rice University (2003), and the Ph.D. from the University of Wisconsin-Madison (2006), all in electrical engineering. He is currently an Assistant Professor at McGill University.