

AN ESCAPE TIME FORMULATION FOR SUBGRAPH DETECTION AND PARTITIONING OF DIRECTED GRAPHS

ZACHARY M. BOYD, NICOLAS FRAIMAN, JEREMY L. MARZUOLA, PETER J. MUCHA,
AND BRAXTON OSTING

ABSTRACT. We provide a rearrangement based algorithm for fast detection of subgraphs of k vertices with long escape times for directed or undirected networks. Complementing other notions of densest subgraphs and graph cuts, our method is based on the mean hitting time required for a random walker to leave a designated set and hit the complement. We provide a new relaxation of this notion of hitting time on a given subgraph and use that relaxation to construct a fast subgraph detection algorithm and a generalization to K -partitioning schemes. Using a modification of the subgraph detector on each component, we propose a graph partitioner that identifies regions where random walks live for comparably large times. Importantly, our method implicitly respects the directed nature of the data for directed graphs while also being applicable to undirected graphs. We apply the partitioning method for community detection to a large class of model and real-world data sets.

1. INTRODUCTION

Subgraph detection and graph partitioning are fundamental problems in network analysis, each typically framed in terms of identifying a group or groups of vertices of the graph so that the vertices in a shared group are well connected or “similar” to each other in their connection patterns while the vertices in different groups (or the complement group) are “dissimilar”. The specific notion of connectedness or similarity is a modeling choice, but one often assumes that edges connect similar vertices, so that in general the detected subgraph is dense and the “communities” identified in graph partitioning are very often more connected within groups than between groups (“assortative communities”).

The identification of subgraphs with particular properties is a long-standing pursuit of network analysis with various applications. Dense subgraphs as assortative communities might represent coordinating regions of interest in the brain [37, 7] or social cliques in a social network [38]. In biology, subgraph detection plays a role in discovering DNA motifs and in gene annotation [25]. In cybersecurity, dense subgraphs might represent anomalous patterns to be highlighted and investigated (e.g., [62]). See [34] for a recent survey and a discussion of alternative computational

DEPARTMENT OF MATHEMATICS, BRIGHAM YOUNG UNIVERSITY, PROVO, UT 84602, USA

DEPARTMENT OF STATISTICS AND OPERATIONS RESEARCH, UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL, CHAPEL HILL, NC 27599, USA

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL, CHAPEL HILL, NC 27599, USA

DEPARTMENT OF MATHEMATICS, DARTMOUTH COLLEGE, HANOVER, NH 03755, USA

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH, SALT LAKE CITY, UT 84112, USA

E-mail addresses: zachboyd@byu.edu, fraiman@email.unc.edu, marzuola@math.unc.edu, peter.j.mucha@dartmouth.edu, osting@math.utah.edu.

The authors are grateful to Matthias Kurzke and Jonathan Weare for helpful discussions throughout the crafting of this result. Z.M.B. and P.J.M. were supported by the ARO under MURI award W911NF-18-1-0244. J.L.M. acknowledges support from the NSF through grant DMS-1909035. P.J.M. was also supported by the NSF through grant BCS-2140024. B. Osting acknowledges support from NSF DMS-1752202. Z.M.B. was supported by DMS-2137511.

methods. As noted there, some of the existing algorithms apply to directed graphs, but most do not.

In the corresponding computer science literature, much of the focus has been on approximation algorithms since the dense k -subgraph is NP-hard to solve exactly (a fact easily seen by a reduction from the k -clique problem). An algorithm that on any input (G, k) returns a subgraph of order k (that is, k vertices or “nodes”; note, we will sometimes refer to the “size” of a graph or subgraph to be the number of vertices, not the number of edges) with average degree within a factor of at most $n^{1/3-\delta}$ from the optimum solution, where n is the order of graph G and $\delta \approx 1/60$ was proposed in [20]. This approximation ratio was the best known for almost a decade until a log-density based approach yielded $n^{1/4+\varepsilon}$ for any $\varepsilon > 0$ [8]. This remains the state-of-the-art approximation algorithm. On the negative side it has been shown [35], assuming the exponential time hypothesis, that there is no polynomial-time algorithm that approximates to within an $n^{1/(\log \log n)^c}$ factor of the optimum. Variations of the problem where the target subgraph has size at most k or at least k have also been considered [1].

Depending on the application of interest, one might seek one or more dense subgraphs within the larger network, a collection of subgraphs to partition the network (i.e., assign a community label to each node), or a set of potentially overlapping subgraphs (see, e.g., [61]). While the literature on “community detection” is enormous (see, e.g., [22, 23, 24, 46, 48] as reviews), a number of common thematic choices have emerged. Many variants of the graph partitioning problem can be formalized as a (possibly constrained) optimization problem. One popular choice minimizes the total weight of the cut edges while making the components roughly equal in size [49]. Another common choice maximizes the total within-community weight relative to that expected at random in some model [42]. Other proposed objective functions include ratio cut weight [11], and approximate “surprise” (improbability) under a cumulative hypergeometric distribution [53]. However, most of these objectives are NP-hard to optimize, leading to the development of a variety of heuristic methods for approximate partitioning (see the reviews cited above for many different approaches). Some of the methods that have been studied are based on the Fiedler eigenvector [21], multicommodity flows [32], semidefinite programming [4, 5, 6], expander flows [3], single commodity flows [29], or Dirichlet partitions [44, 43, 59].

Whichever choice is made for the objective and heuristic, the identified communities can be used to describe the mesoscale structure of the graph and can be important in a variety of applications (see, e.g., the case studies considered in [48]). Subgraphs and communities can also be important inputs to solving problems like graph traversal, finding paths, trees, and flows; while partitioning large networks is often an important sub-problem for complexity reduction or parallel processing in problems such as graph eigenvalue computations [9], breadth-first search [14], triangle listing [17], PageRank [47] and Personalized PageRank [2].

In the present work, we consider a different formulation of the subgraph detection problem, wherein we aim to identify a subgraph with a long mean exit time—that is, the expected time for a random walker to escape the subgraph and hit its complement. Importantly, this formulation inherently respects the possibly directed nature of the edges. This formulation is distinct from either maximizing the total or average edge weight in a dense subgraph and minimizing the edge cut (as a count or suitably normalized) that is necessary to separate a subgraph from its complement. Furthermore, explicitly optimizing for the mean exit time to identify subgraphs may in some applications be preferred as a more natural quantity of interest. For example, in studying the spread of information or a disease on a network, working in terms of exit times is more immediately dynamically relevant than structural measurements of subgraph densities or cuts. Similarly, the development of respondent-driven sampling in the social survey context (see, e.g., [40, 58]) is primarily motivated by there being subpopulations that are difficult to reach (so we expect they often also have high exit times on the directed network with edges reversed). We thus argue that

the identification of subgraphs with large exit times is at least as interesting—and typically related to—those subgraphs with large density and or small cut. Indeed, random walker diffusion on a network and assortative communities are directly related in that the modularity quality function used in many community detection algorithms can be recovered as a low-order truncation of a “Markov stability” auto-correlation measurement of random walks staying in communities [30]. However, the directed nature of the edges is fully respected in our escape time formulation of subgraph detection presented here (cf. random walkers moving either forward or backward along edges in the Markov stability calculation [41] that rederives modularity for a directed network [31]).

From an optimization point of view, the method presented here can be viewed as a rearrangement method or a Merriman-Bence-Osher (MBO) scheme [36] as applied to Poisson solves on a graph. Convergence of MBO schemes is an active area of research in a variety of other scenarios: see [16, 26] in the case of continuum mean curvature flows, [13, 57] in a graph Allen-Cahn type problem, and [27] for a volume constrained MBO scheme on undirected networks. Similarly, proving convergence rates for our algorithm by determining quantitative bounds on the number of interior iterations required for a given ϵ is an important question for the numerical method and its applications to large data sets. Importantly, the method for subgraph detection that we develop and explore, and then extend to a partitioner, is inherently capable of working on directed graphs without any modification. Also, searching for related graph problems where this type of rearrangement algorithm for optimization can be applied will be an important endeavor.

1.1. A New Formulation in Graphs. Let $G = (V, E)$ be a (strongly) connected graph (undirected or directed; we use the term “graph” throughout to include graphs that are possibly directed), with adjacency matrix A with element A_{ij} indicating presence/absence (and possible weight) of an edge from i to j . We define the (out-)degree matrix D to be diagonal with values $D_{ii} = \sum_j A_{ij}$. For weighted edges in A this weighted degree is typically referred to as “strength” but we will continue to use the word “degree” throughout to be this weighted quantity. Consider the discrete time Markov chain M_n for the random walk described by the (row stochastic) *probability transition matrix*, $P := D^{-1}A$. The *exit time from* $S \subset V$ is the stopping time $T_S = \inf\{n \geq 0 : M_n \in S^c\}$. The *mean exit time from* S of a node i is defined by $\mathbb{E}_i T_S$ (where \mathbb{E}_i is the expectation if the walker starts at node i) and is given by v_i , where v is the solution to the system of equations

$$(1.1a) \quad (I - P)_{SS} v_S = 1_S$$

$$(1.1b) \quad v_{S^c} = 0,$$

where the subscript S represents restriction of a vector or matrix to the indices in S . The *average mean escape time (MET) from* S is then

$$(1.2) \quad \tau(S) = \frac{1}{|V|} \sum_{i \in V} v_i,$$

representing the mean exit time from S of a node chosen uniformly at random in the graph (noting that $v_i = 0$ for $i \in S^c$). We are interested in finding vertex sets (of fixed size) having large MET, as these correspond to sets that a random walker would remain in for a long time. Thus, for fixed $k \in \mathbb{N}$, we consider the *subgraph detection problem*,

$$(1.3) \quad \max_{\substack{S \subset V \\ |S|=k}} \tau(S).$$

Multiplying (1.1a) on the left by D , we obtain the equivalent system,

$$(1.4a) \quad Lv = d \text{ on } S,$$

$$(1.4b) \quad v = 0 \text{ on } S^c,$$

where $L = D - A$ is the (unnormalized, out-degree) graph Laplacian, and $d = D\mathbf{1}$ is the out-degree vector. We denote the solution to (1.4) by $v = v(S)$. For $\varepsilon > 0$, we will also consider the approximation to (1.4),

$$(1.5) \quad [L + \varepsilon^{-1}(1 - \phi)] u = d$$

where ϕ is a vector and action by $(1 - \phi)$ on the left is interpreted as multiplication by the diagonal matrix $I - \text{diag}(\phi)$. We denote the solution $u = u_\varepsilon$. Formally, for $\phi = \chi_S$, the characteristic function of S , as $\varepsilon \rightarrow 0$, the vector $u_\varepsilon \rightarrow v_S$ where v_S satisfies (1.4). We can also define an associated approximate MET

$$(1.6) \quad E_\varepsilon(\phi) := \frac{1}{|V|} \|u_\varepsilon\|_{\ell^1(V)} = \frac{1}{|V|} \left\| [L + \varepsilon^{-1}(1 - \phi)]^{-1} d \right\|_{\ell^1(V)},$$

where as $\varepsilon \rightarrow 0$, we have that $E_\varepsilon(\chi_S) \rightarrow \frac{1}{|V|} \|v_S\|_{\ell^1(V)} = \tau(S)$. We then arrive at the following *relaxed subgraph detection problem*

$$(1.7) \quad \max_{\substack{0 \leq \phi \leq 1 \\ \langle \phi, \mathbf{1} \rangle = k}} E_\varepsilon(\phi),$$

which we solve and study in this paper. For small $\varepsilon > 0$, we will study the relationship between the subgraph detection problem (1.3) and its relaxation (1.7).

We are also interested in finding node partitions with high MET in the following sense: Given a vertex subset $S \subset V$, a random walker that starts in S should have difficulty escaping to S^c and a random walker that starts in S^c should have difficulty escaping to S . This leads to the problem $\max_{V = \amalg_{\ell \in [K]} S_\ell} \tau(S) + \tau(S^c)$. More generally, for a vertex partition, $V = \amalg_{\ell \in [K]} S_\ell$, we can consider

$$(1.8) \quad \max_{V = \amalg_{\ell \in [K]} S_\ell} \sum_{\ell \in [K]} \tau(S_\ell).$$

The solution embodies the idea that in a good partition a random walker will transition between partition components very infrequently. An approximation to (1.8) is

$$(1.9) \quad \max_{V = \amalg_{\ell \in [K]} S_\ell} \sum_{\ell \in [K]} E_\varepsilon(\chi_{S_\ell}).$$

We can make an additional approximation by relaxing the constraint set. Define the admissible class

$$\mathcal{A}_K = \left\{ \{\phi_\ell\}_{\ell \in [K]} : \phi_\ell \in \mathbb{R}_+^{|V|} \text{ and } \sum_{\ell \in [K]} \phi_\ell = \mathbf{1} \right\}.$$

Observe that the collection of indicator functions for any K -partition of the vertices is a member of \mathcal{A}_K . Furthermore, we can see that $\mathcal{A}_K \cong (\Delta_K)^{|V|}$, where Δ_K is the unit simplex in K dimensions. Thus, the extremal points of \mathcal{A}_K are precisely the collection of indicator functions for a K -partition of the vertices. For $\varepsilon > 0$, a modified relaxed version of the graph partitioning problem (1.8) can be formulated as

$$(1.10) \quad \min_{\{\phi_\ell\}_{\ell \in [K]} \in \mathcal{A}_K} \tilde{E}_\varepsilon(\{\phi_\ell\}_{\ell \in [K]}), \quad \text{where } \tilde{E}_\varepsilon(\{\phi_\ell\}_{\ell \in [K]}) = \sum_{i=1}^K [1 + \varepsilon |V| E_\varepsilon(\phi_i)]^{-1}.$$

For small $\varepsilon > 0$, we will study the relationship between the graph partitioning problem (1.8) and its relaxation (1.10). An important feature of (1.10) is that it can be optimized using fast rearrangement methods that effectively introduces a volume normalization for the partition sets, while optimization of (1.8) resulted in favoring one partition being full volume. We will discuss this further in Section 2.2.2 below.

1.2. Outline of the Paper. In Section 2, we lay the analytic foundation for rearrangement methods for both the subgraph detection and partitioning problems. We prove the convergence of the methods to local optimizers of our energy functionals in both cases and establish the fact that our fast numerical methods increase the energy. To begin, we establish properties of the gradient and Hessian of the functionals $E_\epsilon(\phi)$ for vectors $0 \leq \phi \leq 1$. Then, using those properties, we introduce rearrangement methods for finding optimizers and prove that our optimization schemes reduce the energy. Then, we discuss how to adapt these results to the partitioning problem. Lastly, we demonstrate how one can easily add a semi-supervised component to our algorithm.

In Section 3, we apply our methods to a variety of model graphs, as well as some empirical data sets to assess their performance. In the subgraph setting, we consider how well we do detecting communities in a family of model graphs related to stochastic block models, made up of a number of random Erdős-Rényi (ER) communities of various sizes and on various scales. The model graphs are designed such that the overall degree distribution is relatively similar throughout. We demonstrate community detectability and algorithm efficacy thresholds by varying a number of parameters in the graph models. We also consider directed graph models of cycles connected to Erdős-Rényi graphs, on which our methods perform quite well. For the partitioners, we also consider related performance studies over our model graph families, as well as on a large variety of clustering data sets.

We conclude in Section 4 with a discussion including possible future directions and applications of these methods.

2. ANALYSIS OF OUR PROPOSED METHODS

In this section, we first analyze the relaxed subgraph detection problem (1.7) and the relaxed graph partitioning problem (1.10). Then, we propose and analyse computational methods for the problems. As noted above, we assume throughout that the graph is (strongly) connected.

2.1. Analysis of the relaxed subgraph detection problem and the relaxed graph partitioning problem. For fixed $\epsilon > 0$ and $\phi \in [0, 1]^{|V|}$, denote the operator on the RHS of (1.5) by $L_\phi := D - A + \frac{1}{\epsilon}(1 - \phi)$.

Lemma 2.1 (Discrete maximum principle). *Given the regularized operator L_ϕ and a vector $f > 0$, we have $(L_\phi^{-1}f)_v > 0$ for all $v \in V$. Without strong connectivity, this result still holds (with $>$ replaced by \geq) as long as there are no leaf nodes.*

Proof. Writing $L_\phi = (D + \frac{1}{\epsilon}(1 - \phi)) - A$, we observe that

$$\begin{aligned} L_\phi^{-1} &= \left(\left(D + \frac{1}{\epsilon}(1 - \phi) \right) \left(I - \left(D + \frac{1}{\epsilon}(1 - \phi) \right)^{-1} A \right) \right)^{-1} \\ &= \left(I - \left(D + \frac{1}{\epsilon}(1 - \phi) \right)^{-1} A \right)^{-1} \left(D + \frac{1}{\epsilon}(1 - \phi) \right)^{-1} \\ &= \sum_{n=0}^{\infty} \left[\left(D + \frac{1}{\epsilon}(1 - \phi) \right)^{-1} A \right]^n \left(D + \frac{1}{\epsilon}(1 - \phi) \right)^{-1}. \end{aligned}$$

Since all entries in the corresponding matrices are positive (by strong connectivity), the result holds. \square

For simplicity, in the following we consider simply setting the potential $X := \epsilon^{-1}(1 - \phi)$ and we use X and $\text{diag } X$ interchangeably for graph Schrödinger operators of the form $L_X := D - A + X$

and solutions of the Poisson equation $L_X u = d$. We can then consider the related energy functional

$$(2.1) \quad E(X) := \left\| [L + X]^{-1} d \right\|_{\ell^1(V)} = \|u\|_{\ell^1(V)}.$$

Lemma 2.2. *The gradient of $E(X)$ with respect to X is given by*

$$(2.2) \quad \nabla E = -u \odot v$$

where \odot denotes the Hadamard product and

$$(2.3) \quad u = (L + X)^{-1} d, \quad v = (L + X)^{-T} e.$$

Here e is the all-ones vector. The Hessian of $E(X)$ with respect to X is then given by

$$(2.4) \quad H = \nabla^2 E = (L + X)^{-1} \odot W + (L + X)^{-T} \odot W^T$$

where

$$W := u \otimes v$$

where \otimes is the Kronecker (or outer) product.

Proof. Write e_j as the indicator vector for the j th entry. First, differentiating (2.3) with respect to X_j , we compute

$$(L + X) \frac{\partial u}{\partial X_j} = -e_j \odot u \quad \implies \quad \frac{\partial u}{\partial X_j} = -\langle e_j, u \rangle (L + X)^{-1} e_j.$$

Taking the second derivative, we obtain

$$\begin{aligned} (L + X) \frac{\partial^2 u}{\partial X_j \partial X_k} &= -e_j \left\langle e_j, \frac{\partial u}{\partial X_k} \right\rangle - e_k \left\langle e_k, \frac{\partial u}{\partial X_j} \right\rangle \\ &= e_j \langle e_k, u \rangle \langle e_j, (L + X)^{-1} e_k \rangle + e_k \langle e_j, u \rangle \langle e_k, (L + X)^{-1} e_j \rangle, \end{aligned}$$

which implies that

$$\frac{\partial^2 u}{\partial X_j \partial X_k} = \langle e_j, (L + X)^{-1} e_k \rangle \langle e_k, u \rangle (L + X)^{-1} e_j + \langle e_k, (L + X)^{-1} e_j \rangle \langle e_j, u \rangle (L + X)^{-1} e_k.$$

By the maximum principle (Lemma 2.1), u is positive and we can write $E(X) = \|u\|_1 = \langle e, u \rangle$. Thus, the gradient is

$$\begin{aligned} \frac{\partial E}{\partial X_j} &= \left\langle e, \frac{\partial u}{\partial X_j} \right\rangle \\ &= -\langle (L + X)^{-T} e, e_j \rangle \langle u, e_j \rangle, \end{aligned}$$

or in other words

$$\nabla_X E = u \odot v$$

for u and v as in (2.3).

For the Hessian, we have

$$\begin{aligned} \frac{\partial^2 E}{\partial X_j \partial X_k} &= \left\langle e, \frac{\partial^2 u}{\partial X_j \partial X_k} \right\rangle \\ &= \langle e_k, (L + X)^{-1} e_j \rangle \langle u, e_j \rangle \langle e_k, v \rangle + \langle e_j, (L + X)^{-1} e_k \rangle \langle v, e_j \rangle \langle e_k, u \rangle. \end{aligned}$$

Thus, the Hessian can be written

$$H = \nabla^2 E = (L + X)^{-1} \odot W + (L + X)^{-T} \odot W^T$$

where

$$W := u \otimes v.$$

as claimed. □

Remark 2.3. If L is symmetric, the above statements can be simplified greatly to give

$$H = \nabla^2 E = (L + X)^{-1} \odot (W + W^T)$$

where

$$W + W^T := u \otimes v + v \otimes u = \frac{1}{2}(u + v) \otimes (u + w) - \frac{1}{2}(u - v) \otimes (u - v).$$

Proposition 2.4. For $f > 0$ fixed, let u satisfy $(L + X)u = f$. The mapping $X \mapsto E(X) = \|u\|_1$ is strongly convex on $\{X \geq 0, X \neq 0\}$.

Proof. We wish to show that

$$E(X) = e^T (L + X)^{-1} d$$

is convex on $[0, X_\infty]^n$ for fixed constant X_∞ . Replacing $D + X$ with X , this is equivalent to

$$e^T (X - A)^{-1} d$$

being convex on $\{V : d_i + X_\infty \geq X_i \geq d_i\}$. Expanding, we have

$$e^T (I - X^{-1}A)^{-1} X^{-1}d = e^T \sum_{k=0}^{\infty} (X^{-1}A)^k X^{-1}d.$$

So it is enough to show that

$$e^T (X^{-1}A)^k X^{-1}d$$

is convex for each $k > 0$. This is true as long as

$$f(x) = \prod_i x_i^{-\alpha_i}$$

is convex for any $\alpha = (\alpha_1, \dots, \alpha_n)$. Computing second derivatives gives

$$f_{X_i X_i}(X) = f(X) \alpha_i (\alpha_i + 1) X_i^{-2}$$

and

$$f_{X_i X_j}(X) = f(X) \alpha_i \alpha_j X_i^{-1} X_j^{-1}.$$

So the Hessian of f is

$$f(X) [(\alpha X^{-1})^T (\alpha X^{-1}) + \text{diag}(\alpha X^{-2})],$$

which is clearly positive semi-definite, being the sum of positive semi-definite matrices.

To observe strong convexity, recognize that the $k = 0$ term contributes a term to the Hessian of the form DX^{-2} , which is positive definite on the domain in question. \square

Proposition 2.4 gives that $\phi \rightarrow E_\varepsilon(\phi)$ is strongly convex on $\mathbb{R}_+^{|V|}$, so $\{\phi_\ell\}_{\ell \in [K]} \mapsto \mathcal{E}^\varepsilon(\{\phi_\ell\}_{\ell \in [K]})$ is also convex on \mathcal{A}_K . The following corollary is then immediate.

Corollary 2.5 (Bang-bang solutions). *Every maximizer of (1.7) is an extreme point of $\{\phi \in [0, 1]^{|V|} : \langle \phi, 1 \rangle = k\}$, i.e., an indicator function for some vertex set $S \subset V$ with $|S| = k$.*

Thus, in the language of control theory, Corollary 2.5 shows that (1.7) is a bang-bang relaxation of (1.3) and that (1.10) is a bang-bang relaxation of (1.8).

Corollary 2.6. *Since the set of values $(x_1, \dots, x_n) \in \mathbb{R}_+^n$ with which we are concerned is convex and E is C^2 in X , the resulting Hessian matrix H is positive definite.*

Remark 2.7. *Note that though the Hadamard product of two positive definite matrices is positive definite, Corollary 2.6 is not obvious from the structure of the Hessian, given that the matrix W is indefinite when u and w are linearly independent. As a result, this positive definiteness is strongly related to the structure of the $L + X$ matrix and its eigenvectors.*

Algorithm 1 Subgraph detector

Input $S^0 \subset V$.

while $S^t \neq S^{t-1}$ **do**

 Solve (2.5) and (2.6) for u and v .

 Assign vertex ℓ to subgraph S^1 if $\nabla_\phi E$ is optimized. That is, solve the following sub-problem.

$$(2.7) \quad \max_{|S|=k} \sum_{\ell \in S} u(\ell) \cdot v(\ell).$$

(Note that (2.7) is easily solved by taking the k indices corresponding to the largest values of $u(\ell) \cdot v(\ell)$, breaking ties randomly if needed.)

 Reset now, building on $S^1 \subset V$ accordingly and repeat until $S^n = S^{n-1}$.

end while

2.2. Optimization scheme.

2.2.1. *Subgraph detector.* We solve (1.7) using rearrangement ideas as follows. After initializing S (randomly in our experiments), we use the gradient (2.2) to find the locally optimal next choice of S , and then iterate until convergence (typically < 10 iterations in our experiments). More explicitly, we follow these steps:

$$(2.5) \quad Lu + \epsilon^{-1}(1 - \chi_{S^0})u = d,$$

$$(2.6) \quad L^T v + \epsilon^{-1}(1 - \chi_{S^0})v = 1.$$

The update, S^1 , then contains those nodes ℓ that maximize $u_\ell v_\ell$.

Pseudocode for this approach is given in Algorithm 1, which has the following ascent guarantee:

Proposition 2.8. *Every nonstationary iteration of algorithm 1 strictly increases the energy E_ϵ . Algorithm 1 terminates in a finite number of iterations.*

Proof. Let S^0 and S^1 be the vertex subsets for successive iterations of the method. Define $W = \chi_{S^1} - \chi_{S^0}$. Assuming $W \neq 0$, by strong convexity (Theorem 2.4) and the formula for the gradient (2.2), we compute

$$(2.8a) \quad E_\epsilon(\chi_{S^1}) > E_\epsilon(\chi_{S^0}) + \frac{1}{\epsilon} \langle W, uv \rangle$$

$$(2.8b) \quad = E_\epsilon(\chi_{S^0}) + \frac{1}{\epsilon} \left(\sum_{i \in S^1} u_i v_i - \sum_{i \in S^0} u_i v_i \right)$$

$$(2.8c) \quad \geq E_\epsilon(\chi_{S^0}).$$

Thus, the energy is strictly increasing on non-stationary iterates. Since we assume that V is a finite size vertex set and the rearrangement method increases the energy, it cannot cycle and hence must terminate in a finite number of iterations. \square

To avoid hand-selection of ϵ , we always set $\epsilon = C/\lambda_F$, where λ_F is the Frobenius norm of the graph Laplacian and $C > 1$ is typically set at $C = 50$ to make sure ϵ allows communication between graph vertices. If C is chosen to take a different value below, we will highlight those cases.

2.2.2. *Graph partitioner.* Given the success of the energy (1.6), one might naively consider partitioning the graph by maximizing an energy of the form

$$(2.9) \quad (S_1, S_2, \dots, S_K) \mapsto \sum_{i=1}^K [E_\epsilon(\chi_{S_i})].$$

However, it can be computed that this energy does not properly constrain the volumes of each partition in a reasonable fashion and the optimizer of this nice problem merely results in putting all the vertices in a single box.

The partition energy we initially worked to minimize instead is of the form

$$(2.10) \quad (S_1, S_2, \dots, S_K) \mapsto \sum_{i=1}^K [|V|E_\epsilon(\chi_{S_i})]^{-1},$$

since the inverses penalize putting all nodes into the same partition by making the resulting empty classes highly costly. Intuitively, this energy functional provides an effective volume normalization of the relative gradients (similar to a K-means type scheme). However, while in practice this functional appears to work reasonably well on all graph models considered here, we were unable to prove, upon analysis of the Hessian, that rearrangements based on such an algorithm are bang-bang like the subgraph detector.

As an alternative, we instead consider

$$(2.11) \quad \tilde{E}_{\delta,\epsilon}(S_1, S_2, \dots, S_K) = \sum_{i=1}^K [1 + \delta|V|E_\epsilon(\chi_{S_i})]^{-1}.$$

Applied to functions, $0 \leq \phi_j \leq 1$, instead of indicator functions, we consider

$$(2.12) \quad \tilde{E}_{\delta,\epsilon}(\phi_1, \phi_2, \dots, \phi_K) = \sum_{i=1}^K [1 + \delta|V|E_\epsilon(\phi_i)]^{-1}.$$

We then have that

$$(2.13) \quad \nabla_{\phi_j} \tilde{E} = -\frac{\delta}{[1 + \delta|V|E_\epsilon(\phi_i)]^2} \nabla_{\phi_j} (|V|E_\epsilon(\phi_j))$$

making the Hessian consist of blocks of the form

$$(2.14) \quad \begin{aligned} \nabla_{\phi_j}^2 \tilde{E} = & -\frac{\delta}{[1 + \delta|V|E_\epsilon(\phi_i)]^2} \nabla_{\phi_j}^2 (|V|E_\epsilon(\phi_j)) \\ & + 2\frac{\delta^2}{[1 + \delta|V|E_\epsilon(\phi_i)]^3} (\nabla_{\phi_j} (|V|E_\epsilon(\phi_j))) (\nabla_{\phi_j} (|V|E_\epsilon(\phi_j)))^T. \end{aligned}$$

Note, this is the sum of a negative definite operator and a rank one matrix, meaning that for δ sufficiently small, the Hessian will prove that \tilde{E} is concave with respect to each component. In practice, we find that taking $\delta = \epsilon$ is sufficient both for having a negative definite Hessian and generating good results with respect to our rearrangement scheme. As such, we will generically take $\delta = \epsilon$ henceforward.

Our approach to the node partitioner is largely analogous to that of the subgraph detector, with the exception that we use class-wise ℓ^1 normalization when comparing which values of $u \cdot v$ at each node. In detail, the algorithm is presented in Algorithm 2. It is a relatively straightforward exercise applying the gradient computation for $E_\epsilon(S_i)$ from Proposition 2.4 to prove that the energy functional (2.10) will decrease with each iteration of our algorithm as in Proposition 2.8.

2.2.3. Semi-supervised learning. In cases where we have a labeled set of nodes T with labels $\hat{\phi}_v \in \{0, 1\}$ indicating whether we want node i to be in the subgraph ($\hat{\phi}_v = 1$) or its complement ($\hat{\phi}_v = 0$), we can incorporate this information into our approach as follows.

For the subgraph detector, we use $E_{\epsilon,\lambda,T}(\phi) = E_\epsilon(\phi) + \lambda \sum_{v \in T} (\phi_v - (1 - \hat{\phi}_v))^2$. Then the rearrangement algorithm needs to be modified at step 3 to become: Assign vertex ℓ to subgraph

Algorithm 2 Graph Partitioner

Input $\vec{S} = \{S_1^0, \dots, S_K^0\}$ a K partition of V .

while $\vec{S}^t \neq \vec{S}^{t-1}$ **do**

For $j = 1, \dots, K$, solve the equations

$$Lu_j + \epsilon^{-1}(1 - \chi_{S_j^0})u_j = d,$$

$$L^T v_j + \epsilon^{-1}(1 - \chi_{S_j^0})v_j = 1.$$

Normalize $u_j = \frac{u_j}{(1 + \epsilon \|u_j\|_{\ell_1})^2}$, $v_j = v_j$.

Assign vertex v to \vec{S}_j^{t+1} where

$$j = \operatorname{argmax}\{u_1 \cdot v_1(v), \dots, u_K \cdot v_K(v)\}$$

(that is, optimize $\nabla_\phi E$) breaking ties randomly if needed.

Set $t = t + 1$.

end while

S^1 if $\nabla_\phi E$ is optimized

$$\max_{|S|=k} \frac{1}{\epsilon} \sum_{\ell \in S} u(\ell) \cdot v(\ell) + 2\lambda \sum_{v \in T} [\chi_S(v) - (1 - \hat{\phi}_v)],$$

where χ is the binary-valued indicator function. This again is solved by picking the largest elements (we break ties by picking the lowest-index maximizers if needed). Since the energy is still convex, the energy still increases at each iteration.

For the K -partitioner, we have a labeled set of nodes T_i with labels $\hat{\phi}_{i,v} \in \{0, 1\}$, for $i = 1, \dots, K$ indicating whether we want node v to be in partition element i , with $\sum_i \hat{\phi}_{i,v} = 1$ for $v \in \cup_i T_i$. We can incorporate this information into our approach by modifying the energy to be the concave functional

$$(2.15) \quad \tilde{E}_{\epsilon, \lambda}(\phi_1, \dots, \phi_K) = \tilde{E}_\epsilon(\phi_1, \dots, \phi_K) - \lambda \sum_{v \in T} \sum_{j=1}^K (\phi_{j,v} - (1 - \hat{\phi}_{j,v}))^2$$

with the gradient rearrangement being appropriately modified.

3. NUMERICAL RESULTS

We test the performance of these algorithms both on synthetic graphs and an assortment of “real-world” graphs. For the synthetic tests, we use a particular set of undirected stochastic block models which we call the MultiScale K -block Escape Ensemble (MICKEE), designed to illustrate some of the data features which our algorithms handle. A MICKEE graph consists of N nodes partitioned into $K + 1$ groups of sizes N_1, \dots, N_K , and $N_{K+1} = N - \sum_{j=1}^K N_j$, where $N_1 < N_2 < \dots < N_K < N_{K+1}$ (see the 2-MICKEE schematic in fig. 1). The nodes in the first K groups induce densely connected Erdős–Rényi (ER) subgraphs (from which we will study escape times) while the last group forms a sparsely connected ER background graph. Each of the K dense subgraphs is sparsely connected to the larger background graph. The goal is to recover one of the planted subgraphs, generally the smallest. A naïve spectral approach will often find one of the planted graphs, but we know of no way to control which subgraph is recovered. Our subgraph detector method, in contrast, can be directed to look at the correct scale to recover a specific subgraph, as we will demonstrate in the 2-MICKEE example (i.e., with two planted subgraphs).




FIGURE 1. Schematic of a 2-MICKEE graph, with three dense subgraphs that are randomly connected to each other. Our subgraph detectors can identify the target subgraph, ignoring other planted subgraphs at different scales. Our partitioner correctly identifies each subgraph as a partition, regardless of the scale.

We explore a number of variations on the basic MICKEE theme, including (1) making the large subgraph have a power law degree distribution (with edges drawn using a loopy, multi-edged configuration model), (2) adding more planted subgraphs with sizes ranging across several scales, (3) adding uniformly random noise edges across the entire graph or specifically between subgraphs, and (4) varying the edge weights of the various types of connections. For brevity, we refer to a MICKEE graph with K planted subgraphs (not including the largest one) as a K -MICKEE graph.

3.1. Subgraph Detection. We explore the performance of Algorithm 1 using four benchmarks, which emphasize (1) noise tolerance, (2) multiscale detection, (3) robustness to heavy-tailed degree distributions, and (4) effective use of directed edges, respectively. In each of these tests, the target subgraph is the smallest planted subgraph.

Robustness to noise. In fig. 2 we visualize results from Algorithm 1 on 3-MICKEE graphs, varying the amount and type of noise. While it is possible to get a bad initialization and thus find a bad local optimum the subgraph detector usually finds the target exactly, except in the noisiest regime (which occurs roughly at the point where the number of noise edges is equal to the number of signal edges).

Range of scales. We generated 2-MICKEE graphs with varying sizes of the subgraphs relative to each other and the total mass. We take $1500 < N < 2500$ for the total size and vary the percentage of smallest planted subgraph as $.02N \leq N_1 \leq .15N$ with $N_2 = 2N_1$. Here, the inter-edge density was set to .01 (in-subgraph-degree values between $(1 - 3p) * N * .01$ for $.02 < p < .15$) with mean inter-edge weight .05 compared to intra-group edge weights of 1. We used this framework to assess the detectability limits of sizes of the smallest components, and numerically we observe that small communities are quite detectable using our algorithm. Using the best result over 5 initializations, we were able to detect the smallest ear over the entire range and we did so reliably on average as well. Since the resulting figure would thus not be terribly informative for this range, we forego including a similar heat plot over this range of parameters.

Heavy-tailed degree distributions. For the results in fig. 3, we use a power law degree distribution in the largest component of 3-MICKEE graphs with $N_1 = 80, N_2 = 160, N_3 = 240$ and $N = 1000$. Surprisingly (at least to us), smaller power-law exponents (corresponding to more skewed degree distributions) actually make the problem much easier (whereas adding noise edges had little effect). We conjecture that this is because, in the presence of very high-degree nodes, it is difficult to have

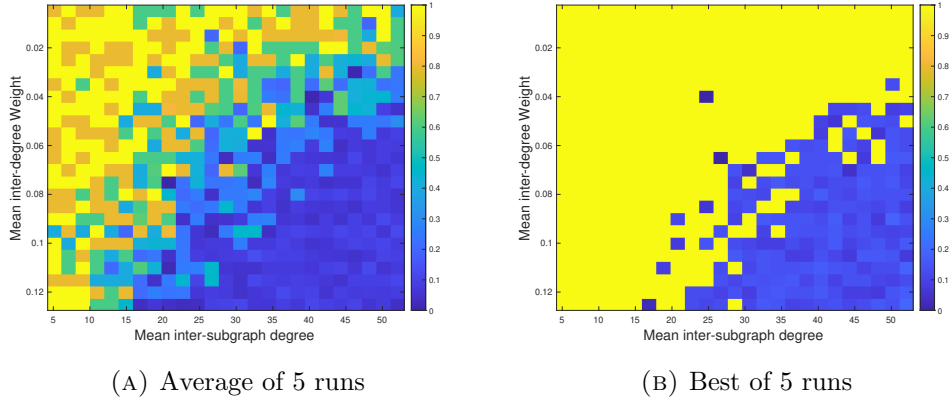


FIGURE 2. Accuracy of Algorithm 1 as a function of mean inter-subgraph degree (the mean taken over the nodes of the target subgraph) and mean weight of the inter-component edges (not including non-edges) for 3-MICKEE graphs with planted subgraphs of sizes 80, 160, and 240 nodes, with a total of 1,000 nodes in the entire graph. The expected in-subgraph-degree is fixed at 20.8 (with intra-component edge weights given by 1). Inter-group edge weights are drawn from a normal distribution with maximum ranging from .01-.25. As long as the noise level is not too high, the subgraph detector finds the smallest planted subgraph despite the presence of “decoy” subgraphs at larger scales. This may be contrasted with spectral clustering, which is attracted to the larger scales.




FIGURE 3. Accuracy of Algorithm 1 on a 3-MICKEE graph with a power law distribution as a function of the power law exponent and inter-cluster edge density. We observe a robustness to both the exponent and density (especially in the right panel) up to a sharp cutoff around 3.4. Note the low exponents (typically considered to be the harder cases) are actually easier in this problem.

a randomly occurring subgraph with high mean escape time, since connections into and out of the hubs are difficult to avoid.

Directed edge utilization. In fig. 4 we consider the problem of detecting a directed cycle appended to an ER graph. The graph weights have been arranged so that the expected degree of all nodes is roughly equal. There are many edges leading from the ER graph into the cycle, with only one edge




FIGURE 4. A directed ER graph with a directed cycle appended. Note that there is only one edge (in the upper left) leading from the cycle to the ER graph, with many edges going the other direction from the ER graph to the cycle. The cycle nodes have the same expected degree as the ER nodes, yet a random walker would naturally get stuck in the cycle for a long time. Detecting such a dynamical trap is a challenge for undirected algorithms, but Algorithm 1 detects it consistently over a wide range of cycle lengths and ER graph sizes.

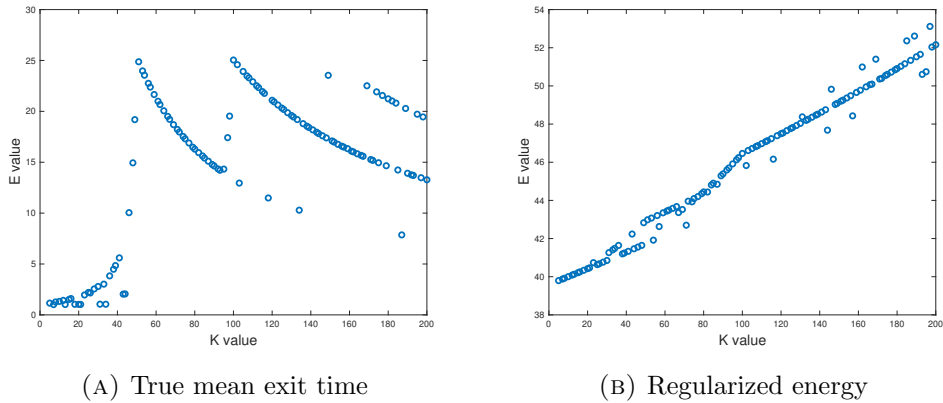


FIGURE 5. The score of the optimal sub-graph found with algorithm 1. Both plots have clear shifts near $k = 50$ corresponding to the smallest component and $k = 100$ corresponding to the second smallest component. This suggests that the size of natural subgraphs within a given graph can be detected from breaks in the subgraph scores as the size of the target in algorithm 1 varies.

leading back into the ER graph. This makes the directed cycle a very salient dynamical feature, but not readily detectable by undirected (e.g. spectral) methods. We considered a large number of cycle sizes relative to the ER graph and with a proper choice of ϵ , we were able to detect the cycle in all cases. Thus, this detector finds directed components very robustly due to the nature of the escape time.

Variation over choice of N_1 . In Figure 5, we consider how the Mean Exit Time as well as the regularized energy in (1.6) behaves as we vary the constrained volume of our algorithm. We considered a 2-MICKEE graph with $N_1 = 50$, $N_2 = 100$ and $N = 1000$. We took the baseline ER density .03 and the inter-edge density was set to .025 with mean inter-edge weight .1.

In summary, we find that the subgraph detector is able to robustly recover planted communities in synthetic graphs and is robust to a range of application-relevant factors.

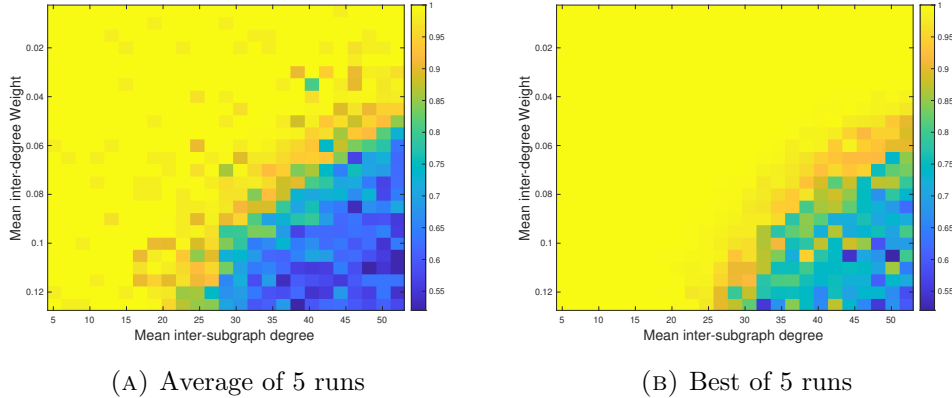


FIGURE 6. The purity measure for algorithm 2 on 3-MICKEE graphs. We vary the density of the inter-region edges and their edge weights. We observe robust (usually perfect) detection over a range of these parameters, with a sharp cutoff (especially in the left panel) when the noise levels grow too high, suggesting that detection is still possible beyond this cutoff, but the energy landscape has more bad local optima beyond this point.

3.2. K -partition method. We will now consider the performance of algorithm 2 in a variety of settings. Throughout, we will give heat plots over the variation of the parameters to visualize the purity measure of our detected communities from our ground-truth smallest component of the graph, over 5 iterations of the algorithm. The purity measure is

$$\frac{1}{N} \sum_{k=1}^K \max_{1 \leq l \leq K} N_k^l$$

for N_k^l the number of data samples in cluster k that are in ground truth class l .

In fig. 6 we consider a $\rho - \Delta$ heat plot of the purity measure for a 4-partition of a 3-MICKEE graph using delocalized connections with $N_1 = 80, N_2 = 160, N_3 = 240$ and $N = 1000$, varying the density of the inter-community edge connections ($0 < \rho < .1$) and the mean weight of the inter-component edges ($0 < \Delta < .125$). We vary over number and strength of connecting edges between components and consider the purity measure as output.

In addition, we have tested algorithm 2 on MICKEE graphs with varying sizes of the components relative to each other and the total mass where the connections between ER graphs include more random edges with weak connection weights. Figure 7 shows results from testing the algorithm on 2-MICKEE graphs with varying sizes of the components relative to each other and the total mass. We take $1500 < N < 2500$ for the total size and vary the percentage of smallest planted subgraph as $.02N \leq N_1 \leq .15N$ with $N_2 = 2N_1$. Here, the inter-edge density was set to .025 with mean inter-edge weight .05. The question addressed in this experiment is how small can we get the components and still detect them. We heat map the average purity measure varying the number of vertices in the graph and the relative size of the smallest sub-graph (i.e., N_1/N).

We similarly consider the partitioning problem on a version of the 3-MICKEE graph with power-law degree distribution in the largest component, using delocalized connections with $N_1 = 80, N_2 = 160, N_3 = 240$ and $N = 1000$. Figure 8 provides a $\rho - q$ plot for results from varying the density ($.001 < \rho < .03$) of the edge-density of connections between the components of the graph, using a power law degree distribution for the largest component with exponent ($2.1 \leq q \leq 4$).

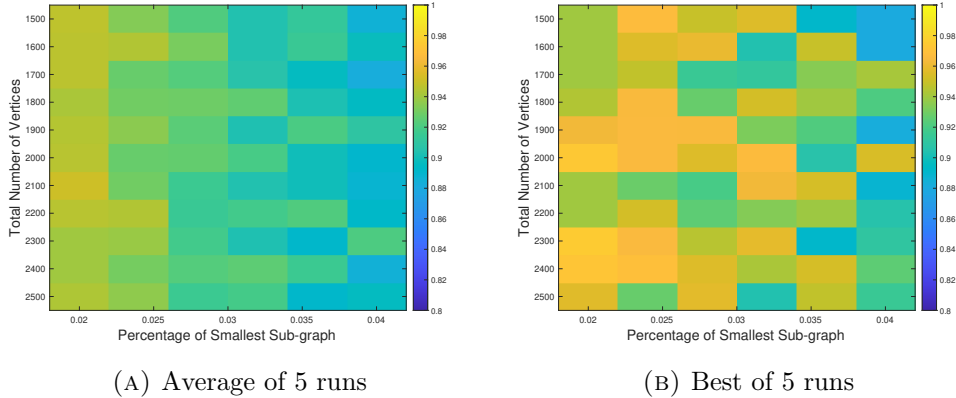


FIGURE 7. The purity measure for the partitioner acting on a 2-MICKEE graph with the fraction of nodes in the smaller planted subgraph varying, along with the size of the graph. We observe a generally robust partitioning.




FIGURE 8. Purity achieved by algorithm 2 on 3-MICKEE graphs with a power law degree distribution, varying the exponent of the power law and inter-subgraph edge density. We observe generally robust partitioning (especially in the right panel).

3.2.1. *Graph clustering examples.* We consider the family of examples as in [63] and compare the best presented purity measures from that paper to a number of settings using our algorithms. Since some of these examples are by their nature actually directed data sets, we throughout computed both the directed and undirected adjacency matrix representations as appropriate to test against. We ran the K -partitioner over a variety of scenarios for both cases. In all these runs, we chose the value of K to agree with the metadata (we avoid term "ground truth", as the node labels themselves may be noisy or not the only good interpretation of the data). However, we note that our algorithm also does a good job in a variety of settings selecting the number of partitions to fill without precisely providing this correct number *a priori*.

For our study, we consider a number of various options for the algorithm. First, the initial seeding sets were chosen either uniformly at random or using K -means on the first K -eigenvectors of the graph Laplacian. We consider the best result over 10 outcomes. In addition, we considered a range of values of ϵ , all of which were a multiplicative factor of the inverse of the Frobenius norm of the graph Laplacian, denoted $\|L\|_{\text{Fro}}$, which sets a natural scaling for separation in the

underlying graph. See for instance the related choice in [44]. We computed a family of partitions for $\epsilon = 50\nu/\|L\|_{\text{Fro}}$, where $\nu = e^{.2\ell}$ with $-50 < \ell < 50$. Finally, we also considered the impact of semi-supervised learning by toggling between $\lambda = 0$ and $\lambda = 10^6$ in (2.15) with 10% of the nodes being included in the learning set. Clearly, there are many ways we might improve the outcomes, by for instance increasing the number and method of initialization and refining our choices of ϵ or λ ; nevertheless, we see under our current choices that our fast algorithm performs well over a range of such parameters, as reported in Table 1.

For each data set in Table 1, we report the best outcome using directed adjacency matrices to build the Graph Laplacian using both the K -means and random initializations but with no semi-supervised learning (Directed); the best outcome using symmetrized adjacency matrices to build the Graph Laplacian using both the K -means and random initializations but with no semi-supervised learning (Undirected); the best outcome when Semi-Supervised Learning is turned on over any configuration (Semi-supervision), the K -means only outcome (K -means only) and the best data from all the experiments reported in [63] (Best from [63]). Our results promisingly demonstrate that our fast algorithm is very successful in many cases in discovering large amounts of community structure that agrees with the metadata in these explicit data sets. Given that our communities are all built around random walks in the graph, it is not clear that all ground-truth designated communities would align well with our methods. For example, we note that our results do not align well with the metadata in the POLBLOGS data set. A major takeaway from the table, however, is that in several examples we see that using the directed nature of the data provides better agreement with the metadata (as indicated by the green cells). Perhaps most striking in the table is that the best run of our fast algorithm, even without semi-supervised learning, provides better agreement with the metadata than [63] for many of the data sets.

As a statistical summary of our findings, we had in total 39 directed datasets and 6 undirected data sets that came from a variety of domains (image, social, biological, physical, etc.). The networks are sized between 35 nodes and 98,528 nodes, having 2-65 classes per network. Among directed networks, 21 data sets gave highest purity with the metadata with semi-supervised learning turned on, while 13 have the best result from [63], and 3 have K -means only best. For 9 total data sets (green in the table), the directed version of our algorithm is better than the symmetrized undirected version, while 5 are tied (yellow) and for 25 the undirected method is better (orange). When [63] is best, the median gap from our result with semi-supervised learning is .05. When our algorithm with semi-supervised learning is best, the median gap from [63] is .05. There is no clear relationship between data domain and performance or node count and performance. However, semi-supervision generally did improve the results the most with a smaller class count (median 3) versus [63] (median 20).

When the directed algorithm is better than the undirected version, the median gap is 0.03. Interestingly, 5 of the datasets where directed was better were image or sensor data, with the two largest gaps (.07 and .11) being digit datasets. When undirected was better, the median gap was 0.06, with the largest gap being .29, for the 20NEWS dataset. When semi-supervision improves over our method (max of directed and undirected performance), the median improvement is .06, and the max improvements were .22 and .20. There is no obvious relationship between edge density and algorithm performance.

We have discussed the output of a variety of experiments on a large number of data sets, but we also want to discuss their dependence upon the ϵ parameter and the percentage of nodes that are learned in the energy (2.15). To that end, we consider the output purity measure for some representative data sets and look at the outputs over a range of epsilon parameters and percentages of learning. In this case, we considered only the K -means initialization for consistency and simplicity of comparison. For the ϵ sweep, we recall that we considered the range $\epsilon = 50\nu/\|L\|_{\text{Fro}}$, where $\nu = e^{.2\ell}$ with $-50 < \ell < 50$. In fig. 9 we show the variation in the purity measure with ϵ for a

Network	Domain	Vertices	Density	Classes	Directed	Undirected	Semi-supervision	K-means only	Best from [63]
Directed data									
MNIST	Digit	70,000	0.00	10	0.85	0.78	0.98	0.84	0.97
VOWEL	Audio	990	0.01	11	0.35	0.32	0.44	0.34	0.37
FAULTS	Materials	1,941	0.00	7	0.44	0.42	0.49	0.39	0.41
SEISMIC	Sensor	98,528	0.00	3	0.60	0.59	0.66	0.58	0.59
7Sectors	Text	4,556	0.00	7	0.27	0.26	0.39	0.26	0.34
PROTEIN	Protein	17,766	0.00	3	0.47	0.46	0.51	0.46	0.50
KHAN	Gene	83	0.06	4	0.59	0.59	0.61	0.59	0.60
ROSETTA	Gene	300	0.02	5	0.78	0.78	0.81	0.77	0.77
WDBC	Medical	683	0.01	2	0.65	0.65	0.70	0.65	0.65
POLBLOGS	Social	1,224	0.01	2	0.55	0.55	0.59	0.51	NA
CITeseer	Citation	3,312	0.00	6	0.28	0.29	0.49	0.25	0.44
SPECT	Astronomy	267	0.02	3	0.79	0.80	0.84	0.79	0.79
DIABETES	Medical	768	0.01	2	0.65	0.67	0.74	0.65	0.65
DUKE	Medical	44	0.11	2	0.64	0.68	0.73	0.52	0.70
IRIS	Biology	150	0.03	3	0.87	0.90	0.97	0.67	0.93
RCV1	Text	9,625	0.00	4	0.35	0.40	0.62	0.32	0.54
CORA	Citation	2,708	0.00	7	0.33	0.39	0.50	0.32	0.47
CURETGREY	Image	5,612	0.00	61	0.23	0.29	0.33	0.22	0.28
SPAM	Email	4,601	0.00	2	0.64	0.70	0.73	0.61	0.69
GISETTE	Digit	7,000	0.00	2	0.87	0.94	0.97	0.81	0.94
WEBKB4	Text	4,196	0.00	4	0.42	0.53	0.66	0.40	0.63
CANCER	Medical	198	0.03	14	0.49	0.55	0.54	0.45	0.54
YALEB	Image	1,292	0.00	38	0.44	0.54	0.52	0.41	0.51
COIL-20	Image	1,440	0.00	20	0.74	0.85	0.78	0.82	0.81
ECOLI	Protein	327	0.02	5	0.79	0.83	0.81	0.81	0.83
YEAST	Biology	1,484	0.00	10	0.46	0.53	0.54	0.47	0.55
20NEWS	Text	19,938	0.00	20	0.20	0.49	0.62	0.16	0.63
MED	Text	1,033	0.00	31	0.50	0.54	0.54	0.48	0.56
REUTERS	Text	8,293	0.00	65	0.60	0.69	0.75	0.60	0.77
ALPHADIGS	Digit	1,404	0.00	6	0.42	0.48	0.48	0.46	0.51
ORL	Face	400	0.01	40	0.76	0.82	0.76	0.78	0.83
OPTDIGIT	Digit	5,620	0.00	10	0.90	0.93	0.91	0.90	0.98
PIE	Face	1,166	0.00	53	0.53	0.66	0.62	0.51	0.74
SEG	Image	2,310	0.00	7	0.54	0.64	0.59	0.51	0.73
UMIST	Face	575	0.01	20	0.74	0.71	0.67	0.67	0.74
PENDIGITS	Digit	10,992	0.00	10	0.82	0.73	0.82	0.83	0.87
SEMEION	Digit	1,593	0.00	10	0.86	0.82	0.77	0.81	0.94
AMLALL	Medical	38	0.13	2	0.92	0.95	0.94	0.95	0.92
IONOSPHERE	Radar	351	0.01	2	0.77	0.77	0.85	0.85	0.70
Undirected data									
POLBOOKS	Social	105	0.08	3	0.83	0.85	0.85	0.82	0.83
KOREA	Social	35	0.11	2	1.00	1.00	1.00	0.71	1.00
FOOTBALL	Sports	115	0.09	12	0.94	0.93	0.90	0.93	0.93
MIREX	Music	3,090	0.00	10	0.21	0.24	0.27	0.12	0.43
HIGHSCHOOL	Social	60	0.10	5	0.82	0.85	0.83	0.82	0.95

TABLE 1. Purity Measure Table

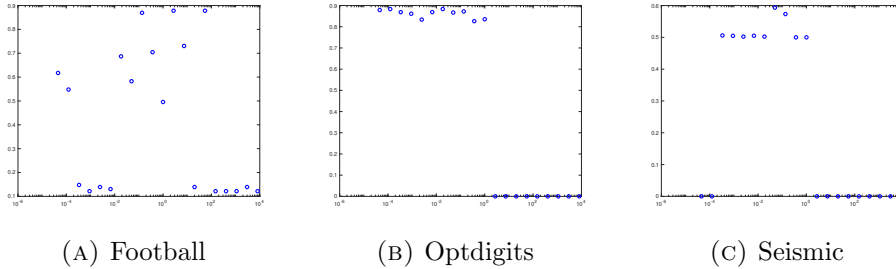


FIGURE 9. Purity measures for three selected data sets as a function of the scale parameter ν . In all three panels, we observe a stable range (on a log scale) where purity is stably nontrivial, and in the left panel, there are two such scales.

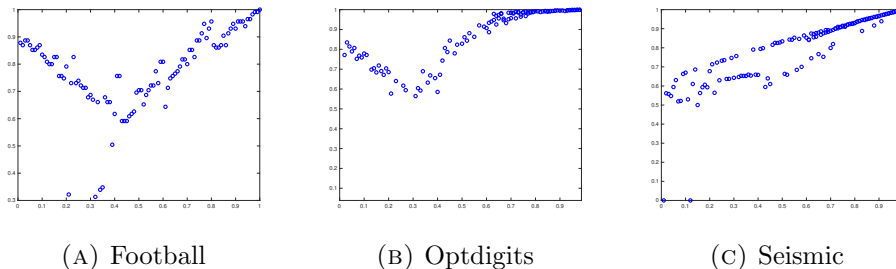


FIGURE 10. Purity measures on three selected data sets as a function of the fraction of supervision (nodes with labels provided) under semi-supervised learning. We observe that supervision can either consistently help (as in the right panel) or can have inconsistent effects (as in the left and middle panels). One possible explanation for this is that there may be multiple clustering structures present in the data, and it takes a lot of supervision to force the partitioner to switch to a partition aligned with the metadata indicated by the supervision, rather than a different clustering structure that is better from the perspective of the optimizer.

small graph (FOOTBALL), a medium sized graph (OPTDIGITS), and a large graph (SEISMIC). Similarly, in fig. 10 we visualize how results vary with the fraction of supervision (nodes with labels provided) under semi-supervised learning, for the same graphs, with $\nu = .6, .8, 1.0, 1.2, 1.4, 1.6, 1.8$.

4. DISCUSSION

Throughout our study we emphasize that our methodology operates fundamentally on the possibly directed nature of the underlying graph data. Considering the Index of Complex Networks [18] as a representative collection of widely-studied networks, we note that (as of our writing here) 327 of the 698 entries in the Index contain directed data. Whereas there are undoubtedly settings where one can ignore edge direction, there are inevitably others where respecting direction is essential. By formulating a strategy for subgraph detection and graph partitioning inherently built on processes running on the directed graph, we avoid the need for any *post hoc* modifications to try to respect directed edges. In particular, our method nowhere relies on any correspondingly undirected version of the graph, avoiding possible information lost in symmetrizing.

While we expect that our formulation of escape times can be useful in general, including for undirected graphs, our proper treatment of the directed graph data should prove especially useful. For example, the directed follower v. following nature of some online social networks (e.g.,

Twitter) is undoubtedly important for understanding the processes involved in the viral spread of (mis)information. As shown by [60] (and extended by [33]), the community structure is particularly important for identifying the virality of memes specifically because a meme that “escapes” (in our present language) its subgraph of origin is typically more likely to continue to propagate. Another application where directed escape times could be relevant is in detecting the (hidden) circulation of information, currency, and resources that is part of coordinated adversarial activity, as explored for example in [28, 39, 51].

To close, we highlight two related thematic areas for possible future work that we believe would lead to important extensions on the methods presented here.

4.1. Connection to distances on directed graphs. In previous work of the present authors [10] along with Jonathan Weare, we construct a symmetrized distance function on the vertices of a directed graph. We recall the details briefly here, which is based somewhat upon the hitting probability matrix construction used in umbrella sampling ([19, 52]). For a general probability transition matrix P , we denote the Perron eigenvector as

$$P'\phi = \phi.$$

Let us define a matrix M such that $M_{ij} = \text{Prob}_i[\tau_j < \tau_i]$, where $\text{Prob}_i[\tau_j < \tau_i]$ is the probability that starting from site i the hitting time of j is less than the time it takes to return to i . Let $X(t)$ be the Markov chain with transition matrix P . Then, it can be observed that ([19, 52])

$$\text{Prob}_i[\tau_j < \tau_i]\phi_i = \text{Prob}_j[\tau_i < \tau_j]\phi_j,$$

where Prob_i represents the probability from $X(0) = i$. This means that from hitting times, one can construct a symmetric adjacency matrix,

$$(4.1) \quad A_{ij}^{(hp)} = \frac{\sqrt{\phi_i}}{\sqrt{\phi_j}} \text{Prob}_i[\tau_j < \tau_i] = A_{ji}^{(hp)}.$$

This adjacency matrix has built-in edge weights based upon hitting times, and we can then easily partition this adjacency matrix using our symmetric algorithms, in particular the mean exit time fast algorithm developed here. The distance function in [10] is given by $d^\beta: [n] \times [n] \rightarrow \mathbb{R}$, which we refer to as the *hitting probability pseudo-metric*, by

$$(4.2) \quad d(i, j) = -\log \left(A_{ij}^{(hp)} \right).$$

This is generically a pseudo-metric as it is possible distinct nodes can be distance 0 from one another, however there exists a quotient graph on which d is a genuine metric. Indeed, a family of metrics is given in [10] that has to do with possible choices of the normalization in (4.1) with different powers of the invariant measure. A natural question to pursue is whether parsing the directed network with this approach to create the symmetrized $A^{(hp)}$ matrix, then applying our clustering scheme can be used to effectively detect graph structures in a more robust manner. In particular, comparison of our clustering scheme versus K -means studies of the distance structure should be an important direction for future study.

4.2. Continuum Limits. The methods presented here have a clear analog in the continuum setting to the motivated problems in the continuum discussed in the introduction. The primary continuum problem is related to the landscape function, or torsion function, on a sub-domain prescribed with Dirichlet boundary conditions,

$$(4.3) \quad -\Delta u_S = 1_S, \quad u_S|_{\partial S} = 0.$$

This is known as the mean exit time from a set S of a standard Brownian motion random walker, see [45], Chapter 7. Correspondingly, for a domain Ω with Neumann boundary conditions (to make life easier with graphs) and some $0 < \alpha < 1$, we propose the following optimization

$$(4.4) \quad \max_{S \subset \Omega, |S| = \alpha |\Omega|} \int_S u_S dx,$$

meaning that we wish to maximize the exit time of a random walker from a given sub-domain. Through the Poisson formula for the mean exit time, we have that $\int u_S = (-\Delta u_S, u_S)$, allowing us to frame things similarly via a Ginzburg–Landau like penalty term for being in a set S ,

$$\min_{\substack{0 \leq \phi \leq 1 \\ \int \phi = \alpha |\Omega|}} \min_{\int u = 1} \frac{1}{2} (-\Delta u_S, u_S) + \frac{1}{2\epsilon} \langle u, (1 - \phi)u \rangle.$$

Analysis of optimizers for such a continuum problem and its use in finding sub-domains and domain partitions is one important direction for future study. Related results in a continuum setting have been studied for instance in [12, 15], but the regularization of this problem seems to be new and connects the problem through the inverse of the Laplacian to the full domain and its boundary conditions. Following works such as [43, 50, 54, 55, 56, 64], an interesting future direction would be to prove consistency of our algorithm to these well-posed continuum optimization problems.

REFERENCES

1. Reid Andersen and Kumar Chellapilla, *Finding dense subgraphs with size bounds*, Int. workshop on algorithms and models for the web-graph, Springer, 2009, pp. 25–37.
2. Reid Andersen, Fan Chung, and Kevin Lang, *Local partitioning for directed graphs using PageRank*, Algorithms and Models for the Web-Graph (Berlin) (Anthony Bonato and Fan R. K. Chung, eds.), Springer, 2007, pp. 166–178.
3. Sanjeev Arora, Elad Hazan, and Satyen Kale, *$O(\log n)$ approximation to SPARSEST CUT in $o(n^2)$ time*, SIAM J. Comp. **39** (2010), no. 5, 1748–1771.
4. Sanjeev Arora, Satish Rao, and Umesh Vazirani, *Expander flows and a $\sqrt{\log n}$ -approximation to sparsest cut*, Proc. 36th ACM Symposium on Theory of Comp., vol. 25, 2004, p. 29.
5. ———, *Geometry, flows, and graph-partitioning algorithms*, Comm. ACM **51** (2008), no. 10, 96–105.
6. ———, *Expander flows, geometric embeddings and graph partitioning*, J. ACM **56** (2009), no. 2, 1–37.
7. Danielle S. Bassett, Nicholas F. Wymbs, Mason A. Porter, Peter J. Mucha, Jean M. Carlson, and Scott T. Grafton, *Dynamic reconfiguration of human brain networks during learning*, Proc. Nat. Acad. Sci. U.S.A. **108** (2011), no. 18, 7641–7646.
8. Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan, *Detecting high log-densities: an $o(n^{1/4})$ approximation for densest k -subgraph*, Proc. 42nd ACM symposium on theory of computing, 2010, pp. 201–210.
9. Erik G. Boman, Karen D. Devine, and Sivasankaran Rajamanickam, *Scalable matrix computations on large scale-free graphs using 2D graph partitioning*, Proc. Int. Conf. High Performance Comput., Networking, Storage Anal., 2013, pp. 1–12.
10. Zachary M. Boyd, Nicolas Fraiman, Jeremy Marzuola, Peter J. Mucha, Braxton Osting, and Jonathan Weare, *A metric on directed graphs and markov chains based on hitting probabilities*, SIAM J. Math. Data Sci. **3** (2021), no. 2, 467–493.
11. Xavier Bresson, Thomas Laurent, David Uminsky, and James H von Brecht, *An adaptive total variation algorithm for computing the balanced cut of a graph*, arXiv:1302.2717, 2013.
12. Tanguy Briançon, *Regularity of optimal shapes for the Dirichlet’s energy with volume constraint*, ESAIM: Control, Optimisation and Calculus of Variations **10** (2004), no. 1, 99–122.
13. Jeremy Budd and Yves Van Gennip, *Graph Merriman–Bence–Osher as a semidiscrete implicit Euler scheme for graph Allen–Cahn flow*, SIAM J. Math. Anal. **52** (2020), no. 5, 4101–4139.
14. Aydın Buluç and Kamesh Madduri, *Graph partitioning for scalable distributed graph computations*, Contemporary Math. **588** (2013), 83–102.
15. Giuseppe Buttazzo and Gianni Dal Maso, *An existence result for a class of shape optimization problems*, Archive Rational Mech. Anal. **122** (1993), no. 2, 183–195.

16. Antonin Chambolle and Matteo Novaga, *Convergence of an algorithm for the anisotropic and crystalline mean curvature flow*, SIAM J. Math. Anal. **37** (2006), no. 6, 1978–1987.
17. Shumo Chu and James Cheng, *Triangle listing in massive networks and its applications*, Proc. 17th ACM Int. Conf. Knowledge Discovery and Data Mining (New York), ACM, 2011, pp. 672–680.
18. Aaron Clauset, Ellen Tucker, and Matthias Sainz, *The Colorado Index of Complex Networks*, 2016, <https://icon.colorado.edu/>.
19. Aaron R. Dinner, Erik H. Thiede, Brian Van Koten, and Jonathan Weare, *Stratification as a general variance reduction method for Markov chain Monte Carlo*, SIAM/ASA J. Uncertainty Quantification **8** (2020), no. 3, 1139–1188.
20. Uriel Feige, David Peleg, and Guy Kortsarz, *The dense k -subgraph problem*, Algorithmica **29** (2001), no. 3, 410–421.
21. Miroslav Fiedler, *Algebraic connectivity of graphs*, Czechoslovak Math. J. **23** (1973), no. 2, 298–305.
22. Santo Fortunato, *Community detection in graphs*, Phys. Rep. **486** (2010), no. 3, 75–174.
23. Santo Fortunato and Darko Hric, *Community detection in networks: A user guide*, Phys. Rep. **659** (2016), 1–44.
24. Santo Fortunato and Mark E. J. Newman, *20 years of network community detection*, Nat. Phys. **18** (2022), no. 88, 848–850 (en).
25. Eugene Fratkin, Brian T Naughton, Douglas L Brutlag, and Serafim Batzoglou, *MotifCut: regulatory motifs finding with maximum density subgraphs*, Bioinformatics **22** (2006), no. 14, e150–e157.
26. Katsuyuki Ishii, *Optimal rate of convergence of the Bence–Merriman–Osher algorithm for motion by mean curvature*, SIAM J. Math. Anal. **37** (2005), no. 3, 841–866.
27. Matt Jacobs, Ekaterina Merkurjev, and Selim Esedoglu, *Auction dynamics: A volume constrained MBO scheme*, J. Comput. Phys. **354** (2018), 288–310.
28. Hui Jin, Xie He, Yanghui Wang, Hao Li, and Andrea L. Bertozzi, *Noisy subgraph isomorphisms on multiplex networks*, 2019 IEEE Int. Conf. on Big Data, 2019, pp. 4899–4905.
29. Rohit Khandekar, Satish Rao, and Umesh Vazirani, *Graph partitioning using single commodity flows*, J. ACM **56** (2009), no. 4, 1–15.
30. Renaud Lambiotte, Jean-Charles Delvenne, and Mauricio Barahona, *Random walks, Markov processes and the multiscale modular organization of complex networks*, IEEE Trans. Net. Sci. Eng. **1** (2014), no. 2, 76–90.
31. E. A. Leicht and M. E. J. Newman, *Community structure in directed networks*, Phys. Rev. Letters **100** (2008), no. 11, 118703–4.
32. Tom Leighton and Satish Rao, *An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms*, Tech. report, MIT Cambridge Microsystems Research Center, 1989.
33. Weihua Li, Skyler J Cranmer, Zhiming Zheng, and Peter J Mucha, *Infectivity enhances prediction of viral cascades in twitter*, PLOS One **14** (2019), no. 4, e0214453.
34. Chenhao Ma, Yixiang Fang, Reynold Cheng, Laks V. S. Lakshmanan, Wenjie Zhang, and Xuemin Lin, *Efficient algorithms for densest subgraph discovery on large directed graphs*, Proc. 2020 ACM SIGMOD Int. Conf. Mgmt. Data (2020), 1051–1066.
35. Pasin Manurangsi, *Almost-polynomial ratio ϵ -hardness of approximating densest k -subgraph*, Proc. 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 954–961.
36. B. Merriman, J. K. Bence, and S. Osher, *Diffusion generated motion by mean curvature*, AMS Selected Letters, Crystal Grower’s Workshop (1993), 73–83.
37. D. Meunier, R. Lambiotte, A. Fornito, K. D. Ersche, and E. T. Bullmore, *Hierarchical modularity in human brain functional networks*, Frontiers in Neuroinformatics **3** (2009), 37.
38. James Moody, *Peer influence groups: Identifying dense clusters in large networks*, Social Networks **23** (2001), no. 4, 261–283.
39. Jacob D. Moorman, Qinyi Chen, Thomas K. Tu, Zachary M. Boyd, and Andrea L. Bertozzi, *Filtering methods for subgraph matching on multiplex networks*, 2018 IEEE Int. Conf. Big Data, 2018, pp. 3980–3985.
40. Ted Mouw and Ashton M. Verdery, *Network sampling with memory: A proposal for more efficient sampling from social networks*, Sociological Methodology **42** (2012), no. 1, 206–256.
41. Peter J. Mucha, Thomas Richardson, Kevin Macon, Mason A. Porter, and Jukka-Pekka Onnela, *Community structure in time-dependent, multiscale, and multiplex networks*, Science **328** (2010), no. 5980, 876–878.
42. M. E. J. Newman and M. Girvan, *Finding and evaluating community structure in networks*, Phys. Rev. E **69** (2004), no. 2, 026113.
43. Braxton Osting and Todd Harry Reeb, *Consistency of Dirichlet partitions*, SIAM J. Math. Anal. **49** (2017), no. 5, 4251–4274.
44. Braxton Osting, Chris D White, and Édouard Oudet, *Minimal Dirichlet energy partitions for graphs*, SIAM J. Sci. Comp. **36** (2014), no. 4, A1635–A1651.

45. Grigorios A Pavliotis, *Stochastic processes and applications: Diffusion processes, the Fokker–Planck and Langevin equations*, vol. 60, Springer, 2014.
46. M. A. Porter, J. P. Onnela, and P. J. Mucha, *Communities in networks*, Notices of the AMS **56** (2009), no. 9, 1082–1097 & 1164–1166.
47. Semih Salihoglu and Jennifer Widom, *GPS: A graph processing system*, SSDBM: Proc. 25th Int. Conf. Scientific and Stat. Database Mgmt. (New York), SSDBM, ACM, 2013.
48. Saray Shai, Natalie Stanley, Clara Granell, Dane Taylor, and Peter J. Mucha, *Case studies in network community detection*, The Oxford Handbook of Social Networks (Ryan Light and James Moody, eds.), Oxford University Press, 2021, pp. 309–333.
49. Jianbo Shi and Jitendra Malik, *Normalized cuts and image segmentation*, IEEE Trans. Pattern Anal. Machine Intelligence **22** (2000), no. 8, 888–905.
50. Amit Singer and Hau-Tieng Wu, *Spectral convergence of the connection Laplacian from random samples*, Information and Inference **6** (2017), no. 1, 58–123.
51. Daniel L. Sussman, Youngser Park, Carey E. Priebe, and Vince Lyzinski, *Matched filters for noisy induced subgraph detection*, IEEE Trans. Pattern Analysis and Machine Intelligence **42** (2020), no. 11, 2887–2900.
52. Erik Thiede, Brian Van Koten, and Jonathan Weare, *Sharp entrywise perturbation bounds for Markov chains*, SIAM J. Matrix Anal. App. **36** (2015), no. 3, 917–941.
53. V. A. Traag, R. Aldecoa, and J.-C. Delvenne, *Detecting communities using asymptotical surprise*, Phys. Rev. E **92** (2015), no. 2, 022816.
54. Nicolás García Trillos and Dejan Slepčev, *Continuum limit of total variation on point clouds*, Archive for Rational Mechanics and Analysis **220** (2016), no. 1, 193–241.
55. ———, *A variational approach to the consistency of spectral clustering*, Appl. Comp. Harmonic Anal. **45** (2018), no. 2, 239–281.
56. Nicolás García Trillos, Dejan Slepčev, James Von Brecht, Thomas Laurent, and Xavier Bresson, *Consistency of Cheeger and ratio graph cuts*, J. Machine Learning Research **17** (2016), no. 1, 6268–6313.
57. Yves van Gennip, Nestor Guillen, Braxton Osting, and Andrea L. Bertozzi, *Mean curvature, threshold dynamics, and phase field theory on finite graphs*, Milan J. Math. **82** (2014), no. 1, 3–65.
58. Ashton M. Verdery, M. Giovanna Merli, James Moody, Jeffrey A. Smith, and Jacob C. Fisher, *Brief report: Respondent-driven sampling estimators under real and theoretical recruitment conditions of female sex workers in China*, Epidemiology **26** (2015), no. 5, 661–665.
59. Dong Wang and Braxton Osting, *A diffusion generated method for computing Dirichlet partitions*, J. Comp. Appl. Math. **351** (2019), 302–316.
60. Lilian Weng, Filippo Menczer, and Yong-Yeol Ahn, *Virality prediction and community structure in social networks*, Sci. Rep. **3** (2013), 2522.
61. James D. Wilson, Simi Wang, Peter J. Mucha, Shankar Bhamidi, and Andrew B. Nobel, *A testing based extraction algorithm for identifying significant communities in networks*, Ann. Appl. Stat. **8** (2014), no. 3, 1853–1891.
62. Hao Yan, Qianzhen Zhang, Deming Mao, Ziyue Lu, Deke Guo, and Sheng Chen, *Anomaly detection of network streams via dense subgraph discovery*, 2021 Int. Conf. Comput. Communications and Networks, Jul 2021, pp. 1–9.
63. Zhirong Yang, Tele Hao, Onur Dikmen, Xi Chen, and Erkki Oja, *Clustering by nonnegative matrix factorization using graph random walk*, Adv. Neural Inf. Proc. Sys. **25** (2012), 1079–1087.
64. A. Yuan, J. Calder, and B. Osting, *A continuum limit for the PageRank algorithm*, European J. Appl. Math. **33** (2021), no. 3, 472–504.