

A Generalized Matching Pursuit Approach for Graph-Structured Sparsity

Feng Chen, Baojian Zhou

Computer Science Department, University at Albany – SUNY
 1400 Washington Avenue, Albany, NY, USA
 {fchen5, bzhou6}@albany.edu

Abstract

Sparsity-constrained optimization is an important and challenging problem that has wide applicability in data mining, machine learning, and statistics. In this paper, we focus on sparsity-constrained optimization in cases where the cost function is a general nonlinear function and, in particular, the sparsity constraint is defined by a **graph-structured sparsity** model. Existing methods explore this problem in the context of sparse estimation in linear models. To the best of our knowledge, this is the first work to present an efficient approximation algorithm, namely, GRAPH-structured Matching Pursuit (GRAPH-MP), to optimize a general nonlinear function subject to graph-structured constraints. We prove that our algorithm enjoys the strong guarantees analogous to those designed for linear models in terms of convergence rate and approximation accuracy. As a case study, we specialize GRAPH-MP to optimize a number of well-known graph scan statistic models for the connected subgraph detection task, and empirical evidence demonstrates that our general algorithm performs superior over state-of-the-art methods that are designed specifically for the task of connected subgraph detection.

1 Introduction

In recent years, that is a growing demand on efficient computational methods for analyzing high-dimensional data in a variety of applications such as bioinformatics, medical imaging, social networks, and astronomy. In many settings, sparsity has been shown effective to model latent structure in high-dimensional data and at the same time remain a mathematically tractable concept. Beyond the ordinary, extensively studied, sparsity model, a variety of **structured sparsity models** have been proposed in the literature, such as the sparsity models defined through trees [Hegde *et al.*, 2014b], groups [Jacob *et al.*, 2009], clusters [Huang *et al.*, 2011], paths [Asteris *et al.*, 2015], and connected subgraphs [Hegde *et al.*, 2015b]. These sparsity models are designed to capture the interdependence of the locations of the non-zero components via prior knowledge, and are considered in the general

sparsity-constrained optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } \text{supp}(\mathbf{x}) \in \mathbb{M}, \quad (1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a differentiable cost function and the sparsity model \mathbb{M} is defined as a family of structured supports: $\mathbb{M} = \{S_1, S_2, \dots, S_L\}$, where $S_i \subseteq [n]$ satisfies a certain structure property (e.g., trees, groups, clusters). The original k -sparse recovery problem corresponds to the particular case where the model $\mathbb{M} = \{S \subseteq [n] \mid |S| \leq k\}$.

The methods that focus on general nonlinear cost functions fall into two major categories, including **structured sparsity-inducing norms based** and **model-projection based**, both of which often assume that the cost function $f(\mathbf{x})$ satisfies a certain convexity/smoothness condition, such as *Restricted Strong Convexity/Smoothness* (RSC/RSS) or *Stable Mode-Restricted Hessian* (SMRH). In particular, the methods in the first category replace the structured sparsity model with regularizations by a sparsity-inducing norm that is typically non-smooth and non-Euclidean [Bach *et al.*, 2012]. The methods in the second category decompose Problem (1) into an unconstrained subproblem and a model projection oracle that finds the best approximation of an arbitrary \mathbf{x} in the model \mathbb{M} :

$$P(\mathbf{x}) = \arg \min_{\mathbf{x}' \in \mathbb{M}} \|\mathbf{x} - \mathbf{x}'\|_2^2 \text{ s.t. } \text{supp}(\mathbf{x}') \in \mathbb{M}.$$

A number of methods are proposed specifically for the k -sparsity model $\mathbb{M} = \{S \subseteq [n] \mid |S| \leq k\}$, including the forward-backward algorithm [Zhang, 2009], the gradient descent algorithm [Tewari *et al.*, 2011], the gradient hard-thresholding algorithms [Yuan *et al.*, 2013; Bahmani *et al.*, 2013; Jain *et al.*, 2014], and the Newton greedy pursuit algorithm [Yuan and Liu, 2014]. A limited number of methods are proposed for other types of structured sparsity models via projected gradient descent, such as the union of subspaces [Blumensath, 2013] and the union of nested subsets [Bahmani *et al.*, 2016].

In this paper, we focus on general nonlinear optimization subject to graph-structured sparsity constraints. Our approach applies to data with an underlying graph structure in which nodes corresponding to $\text{supp}(\mathbf{x})$ form a small number of connected components. By a proper choice of the underlying graph, several other structured sparsity models such as the “standard” k -sparsity, block sparsity, cluster sparsity, and tree sparsity can be encoded as special cases of graph-structured sparsity [Hegde *et al.*, 2015a].

We have two key observations: 1) **Sparsity-inducing norms.** There is no known sparsity-inducing norm that is able to capture graph-structured sparsity. The most relevant norm is generalized fused lasso [Xin *et al.*, 2014] that enforces the smoothness between neighboring entries in \mathbf{x} , but does not have fine-grained control over the number of connected components. Hence, existing methods based on sparsity-inducing norms are not directly applicable to the problem to be optimized. 2) **Model projection oracle.** There is no exact model projection oracle for a graph-structured sparsity model, as this exact projection problem is NP-hard due to a reduction from the classical Steiner tree problem [Hegde *et al.*, 2015b]. As most existing model-projection based methods assume an exact model projection oracle, they are not directly applicable here as well. To the best of our knowledge, there is only one recent approach that admits inexact projections for a graph-structured sparsity model by assuming “head” and “tail” approximations for the projections, but is only applicable to linear regression problems [Hegde *et al.*, 2015b]. This paper will generalize this approach to optimize general nonlinear functions. The main contributions of our study are summarized as follows:

- **Design of an efficient approximation algorithm.** A new and efficient algorithm, namely, GRAPH-MP, is developed to approximately solve Problem (1) with a differentiable cost function and a graph-structured sparsity model. We show that GRAPH-MP reduces to a state-of-the-art algorithm for graph-structured compressive sensing and linear models, namely, GRAPH-COSAMP, when $f(\mathbf{x})$ is a least square loss function.
- **Theoretical analysis and connections to existing methods.** The convergence rate and accuracy of the proposed GRAPH-MP are analyzed under a condition of $f(\mathbf{x})$ that is weaker than popular conditions such as RSC/RSS and SMRH. We demonstrate that GRAPH-MP enjoy strong guarantees analogous to GRAPH-COSAMP on both convergence rate and accuracy.
- **Compressive experiments to validate the effectiveness and efficiency of the proposed techniques.** The proposed GRAPH-MP is applied to optimize a variety of graph scan statistic models for the task of connected subgraph detection. Extensive experiments demonstrate that GRAPH-MP performs superior over state-of-the-art methods that are customized for the task of connected subgraph detection on both running time and accuracy.

The rest of this paper is organized as follows. Section 2 introduces the graph-structured sparsity model. Section 3 formalizes the problem and presents an efficient algorithm GRAPH-MP. Sections 4 and 5 present theoretical analysis and applications of GRAPH-MP. Experiments are presented in Section 6, and Section 7 describes future work. An extended version of this paper is available from [Chen, 2016].

2 Graph-Structured Sparsity Model

Given an underlying graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ defined on the coefficients of the unknown vector \mathbf{x} , where $\mathbb{V} = [n]$ and $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$, a graph-structured sparsity model has the form:

$$\mathbb{M}(k, g) = \{S \subseteq \mathbb{V} \mid |S| \leq k, \gamma(S) = g\}, \quad (2)$$

where k refers to an upper bound of the sparsity (total number of nodes) of S and $\gamma(S) = g$ refers to the maximum number of connected components formed by the forest induced by S : $\mathbb{G}_S = (S, \mathbb{E}_S)$, where $\mathbb{E}_S = \{(i, j) \mid i, j \in S, (i, j) \in \mathbb{E}\}$. The corresponding model projection oracle is defined as

$$P(\mathbf{x}) = \arg \min_{\mathbf{x}' \in \mathbb{R}^n} \|\mathbf{x} - \mathbf{x}'\|_2^2 \text{ s.t. } \text{supp}(\mathbf{x}') \in \mathbb{M}(k, g). \quad (3)$$

Solving Problem (3) exactly is NP-hard due to a reduction from the classical Steiner tree problem. Instead of solving (3) exactly, two nearly-linear time approximation algorithms with the following complementary approximation guarantees are proposed in [Hegde *et al.*, 2015b]:

- **Tail approximation (T(x)):** Find $S \in \mathbb{M}(k_T, g)$ such that
$$\|\mathbf{x} - \mathbf{x}_S\|_2 \leq c_T \cdot \min_{S' \in \mathbb{M}(k, g)} \|\mathbf{x} - \mathbf{x}_{S'}\|_2, \quad (4)$$

where $c_T = \sqrt{7}$ and $k_T = 5k$.

- **Head approximation (H(x)):** Find $S \in \mathbb{M}(k_H, g)$ such that
$$\|\mathbf{x}_S\|_2 \geq c_H \cdot \max_{S' \in \mathbb{M}(k, g)} \|\mathbf{x}_{S'}\|_2, \quad (5)$$

where $c_H = \sqrt{1/14}$ and $k_H = 2k$.

If $c_T = c_H = 1$, then $T(\mathbf{x}) = H(\mathbf{x}) = S$ provides the exact solution of the model projection oracle: $P(\mathbf{x}) = \mathbf{x}_S$, which indicates that the approximations stem from the fact that $c_T > 1$ and $c_H < 1$. We note that these two approximations originally involve additional budgets (B) based on edge weights, which are ignored in this paper by setting unit edge weights and $B = k - g$.

Generalization: The above graph-structured sparsity model is defined based on the number of connected components in the forest induced by S . This model can be generalized to graph-structured sparsity models that are defined based on other graph topology constraints, such as density, k -core, radius, cut, and various others, as long as their corresponding head and tail approximations are available [Chen, 2016].

3 Problem Statement and Algorithm

Given the graph-structured sparsity model, $\mathbb{M}(k, g)$, as defined above, the sparsity-constrained optimization problem to be studied is formulated as:

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } \text{supp}(\mathbf{x}) \in \mathbb{M}(k, g), \quad (6)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a differentiable cost function, and the upper bound of sparsity k and the maximum number of connected components g are predefined by users.

Hegde *et al.* propose GRAPH-COSAMP, a variant of COSAMP [Hegde *et al.*, 2015b] to optimize the least square cost function $f(\mathbf{x}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$ based on the head and tail approximations. The authors show that GRAPH-COSAMP achieves an information-theoretically optimal sample complexity for a wide range of parameters. In this paper, we generalize GRAPH-COSAMP and propose a new algorithm named as GRAPH-MP for Problem (6), as shown in Algorithm 1. The first step (Line 3) in each iteration, $\mathbf{g} = \nabla f(\mathbf{x}^i)$, evaluates the gradient of the cost function at the current estimate. Then a subset of nodes are identified via head approximation, $\Gamma = H(\mathbf{g})$, that returns a support set with

head value at least a constant fraction of the optimal head value, in which pursuing the minimization will be most effective. This subset is then merged with the support of the current estimate to obtain the merged subset Ω , over which the function f is minimized to produce an intermediate estimate, $\mathbf{b} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$ s.t. $\mathbf{x}_{\Omega^c} = 0$. Then a subset of nodes are identified via tail approximation, $B = \mathbf{T}(\mathbf{b})$, that returns a support set with tail value at most a constant times larger than the optimal tail value. The iterations terminate when the halting condition holds. There are two popular options to define the halting condition: 1) the change of the cost function from the previous iteration is less than a threshold ($|f(\mathbf{x}^{i+1}) - f(\mathbf{x}^i)| \leq \epsilon$); and 2) the change of the estimated minimum from the previous iteration is less than a threshold ($\|\mathbf{x}^{i+1} - \mathbf{x}^i\|_2 \leq \epsilon$), where ϵ is a predefined threshold (e.g., $\epsilon = 0.001$).

Algorithm 1 GRAPH-MP

```

1:  $i = 0, \mathbf{x}^i = 0$ ;
2: repeat
3:    $\mathbf{g} = \nabla f(\mathbf{x}^i)$ ;
4:    $\Gamma = \mathbf{H}(\nabla f(\mathbf{g}^i))$ ;
5:    $\Omega = \Gamma \cup \text{supp}(\mathbf{x}^i)$ 
6:    $\mathbf{b} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$  s.t.  $\mathbf{x}_{\Omega^c} = 0$ 
7:    $B = \mathbf{T}(\mathbf{b})$ ;
8:    $\mathbf{x}^{i+1} = \mathbf{b}_B$ 
9: until halting condition holds
10: return  $\mathbf{x}^{i+1}$ 

```

4 Theoretical Analysis of GRAPH-MP

In order to demonstrate the accuracy of estimates using Algorithm 1 we require a variant of the *Restricted Strong Convexity/Smoothness* (RSC/RSS) conditions proposed in [Yuan et al., 2013] to hold. The RSC condition basically characterizes cost functions that have quadratic bounds on the derivative of the objective function when restricted to model-sparse vectors. The condition we rely on, the Weak Restricted Strong Convexity (WRSC), can be formally defined as follows:

Definition 1 (Weak Restricted Strong Convexity Property (WRSC)). *A function $f(\mathbf{x})$ has condition $(\xi, \delta, \mathbb{M})$ -WRSC if $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\forall S \in \mathbb{M}$ with $\text{supp}(\mathbf{x}) \cup \text{supp}(\mathbf{y}) \subseteq S$, the following inequality holds for some $\xi > 0$ and $0 < \delta < 1$:*

$$\|\mathbf{x} - \mathbf{y} - \xi \nabla_S f(\mathbf{x}) + \xi \nabla_S f(\mathbf{y})\|_2 \leq \delta \|\mathbf{x} - \mathbf{y}\|_2. \quad (7)$$

Remark 1. 1) In the special case where $f(\mathbf{x}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$ and $\xi = 1$, condition $(\xi, \delta, \mathbb{M})$ -WRSC reduces to the well known *Restricted Isometry Property (RIP)* condition in compressive sensing. 2) The RSC and RSS conditions imply condition WRSC, which indicates that condition WRSC is no stronger than the RSC and RSS conditions [Yuan et al., 2013].

Theorem 4.1. *Consider the graph-structured sparsity model $\mathbb{M}(k, g)$ for some $k, g \in \mathbb{N}$ and a cost function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that satisfies condition $(\xi, \delta, \mathbb{M}(8k, g))$ -WRSC. If $\eta = c_H(1 - \delta) - \delta > 0$, then for any true $\mathbf{x} \in \mathbb{R}^n$ with $\text{supp}(\mathbf{x}) \in \mathbb{M}(k, g)$, the iterates of Algorithm 1 obey*

$$\|\mathbf{x}^{i+1} - \mathbf{x}\|_2 \leq \alpha \|\mathbf{x}^i - \mathbf{x}\|_2 + \beta \|\nabla_I f(\mathbf{x})\|, \quad (8)$$

$$\text{where } \beta = \frac{\xi(1+c_T)}{1-\delta} \left[\frac{(1+c_H)}{\eta} + \frac{\eta(1+c_H)}{\sqrt{1-\eta^2}} + 1 \right], \quad \alpha = \frac{(1+c_T)}{1-\delta} \sqrt{1-\eta^2}, \text{ and } I = \arg \max_{S \in \mathbb{M}(8k, g)} \|\nabla_S f(\mathbf{x})\|_2.$$

Proof. A proof of this result can be found in Appendix. \square

As indicated in Theorem 4.1, under proper conditions the estimator error of GRAPH-MP is determined by the multiplier of $\|\nabla_S f(\mathbf{x})\|_2$, and the convergence rate before reaching this error level is geometric. In particular, if the true \mathbf{x} is sufficiently close to an unconstrained minimum of f , then the estimation error is negligible because $\|\nabla_S f(\mathbf{x})\|_2$ has a small magnitude. Especially, in the ideal case where $\nabla f(\mathbf{x}) = 0$, it is guaranteed that we can obtain the true \mathbf{x} to arbitrary precision. If we further assume that $\alpha = \frac{(1+c_T)\sqrt{1-\eta^2}}{\sqrt{1-\delta}} < 1$, then exact recovery can be achieved in finite iterations.

The shrinkage rate $\alpha < 1$ controls the convergence of GRAPH-MP, and it implies that when δ is very small, the approximation factors c_H and c_T satisfy

$$c_H^2 > 1 - 1/(1 + c_T)^2. \quad (9)$$

We note that the head and tail approximation algorithms designed in [Hegde et al., 2015b] do not satisfy the above condition, with $c_T = \sqrt{7}$ and $c_H = \sqrt{1/14}$. However, as proved in [Hegde et al., 2015b], the approximation factor c_H of any given head approximation algorithm can be **boosted** to any arbitrary constant $c'_H < 1$, such that the above condition is satisfied. Empirically it is not necessary to boost the head-approximation algorithm as strongly as suggested by the analysis in [Hegde et al., 2014a].

Theorem 4.2. *Let $\mathbf{x} \in \mathbb{R}^n$ be a true optimum such that $\text{supp}(\mathbf{x}) \in \mathbb{M}(k, g)$, and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a cost function that satisfies condition $(\xi, \delta, \mathbb{M}(8k, g))$ -WRSC. Assuming that $\alpha < 1$, GRAPH-MP returns a $\hat{\mathbf{x}}$ such that, $\text{supp}(\hat{\mathbf{x}}) \in \mathbb{M}(5k, g)$ and $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq c \|\nabla_I f(\mathbf{x})\|_2$, where $c = (1 + \frac{\beta}{1-\alpha})$ is a fixed constant. Moreover, GRAPH-MP runs in time*

$$O((T + |\mathbb{E}| \log^3 n) \log(\|\mathbf{x}\|_2 / \|\nabla_I f(\mathbf{x})\|_2)), \quad (10)$$

where T is the time complexity of one execution of the subproblem in Line 6. In particular, if T scales linearly with n , then GRAPH-MP scales nearly linearly with n .

Proof. The i -th iterate of Algorithm 1 satisfies

$$\|\mathbf{x} - \mathbf{x}^i\|_2 \leq \alpha^i \|\mathbf{x}\|_2 + \frac{\beta}{1-\alpha} \|\nabla_I f(\mathbf{x})\|_2. \quad (11)$$

After $t = \left\lceil \log \left(\frac{\|\mathbf{x}\|_2}{\|\nabla_I f(\mathbf{x})\|_2} \right) / \log \frac{1}{\alpha} \right\rceil$ iterations, Algorithm 1 returns an estimate $\hat{\mathbf{x}}$ satisfying $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq (1 + \frac{\beta}{1-\alpha}) \|\nabla_I f(\mathbf{x})\|_2$. The time complexities of both head and tail approximations are $O(|\mathbb{E}| \log^3 n)$. The time complexity of one iteration in Algorithm 1 is $(T + |\mathbb{E}| \log^3 n)$, and the total number of iterations is $\left\lceil \log \left(\frac{\|\mathbf{x}\|_2}{\|\nabla_I f(\mathbf{x})\|_2} \right) / \log \frac{1}{\alpha} \right\rceil$, and the overall time complexity follows. \square

Remark 2. *The previous algorithm GRAPH-COSAMP [Hegde et al., 2015b] for compressive sensing is a special case of GRAPH-MP. Assume $f(\mathbf{x}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$.*

1) **Reduction.** The gradient in Step 3 of GRAPH-MP has the form: $\nabla f(\mathbf{x}^i) = -\mathbf{A}^T(\mathbf{y} - \mathbf{A}\mathbf{x}^i)$, and an analytical form of \mathbf{b} in Step 6 can be obtained as: $\mathbf{b}_\Omega = \mathbf{A}_\Omega^+ \mathbf{y}$ and $\mathbf{b}_{\Omega^c} = 0$, where $\mathbf{A}^+ = \mathbf{A}^T(\mathbf{A}^T\mathbf{A})^{-1}$, which indicates that GRAPH-MP reduces to GRAPH-COSAMP in this scenario. 2) **Shrinkage rate.** The shrinkage rate α of GRAPH-MP is analogous to that of GRAPH-COSAMP, even though that the shrinkage rate of GRAPH-COSAMP is optimized based on the RIP sufficient constants. In particular, they are identical when δ is very small. 3) **Constant component.** Assume that $\xi = 1$. Condition $(\xi, \delta, \mathbb{M}(k, g))$ -WRSC then reduces to the RIP condition in compressive sensing. Let $\mathbf{e} = \mathbf{y} - \mathbf{A}\mathbf{x}$. The component $\|\nabla f(\mathbf{x}^i)\|_2 = \|\mathbf{A}^T \mathbf{e}\|_2$ is upper bounded by $\sqrt{1 + \delta} \|\mathbf{e}\|_2$ [Hegde et al., 2014a]. The constant $\beta \|\nabla_1 f(\mathbf{x})\|$ is then upper bounded by $\frac{\xi(1+c_T)\sqrt{1+\delta}}{1-\delta} \left[\frac{(1+c_H)}{\eta} + \frac{\eta(1+c_H)}{\sqrt{1-\eta^2}} + 1 \right] \|\mathbf{e}\|_2$ that is analogous to the constant of GRAPH-COSAMP, and they are identical when δ is very small.

5 Application in Graph Scan Statistic Models

In this section, we specialize GRAPH-MP to optimize a number of graph scan statistic models for the task of connected subgraph detection. Given a graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, where $\mathbb{V} = [n]$, $\mathbb{E} \subseteq \mathbb{V} \times \mathbb{V}$, and each node v is associated with a vector of features $\mathbf{c}(v) \in \mathbb{R}^p$. Let $S \subseteq \mathbb{V}$ be a connected subset of nodes. A graph scan statistic, $F(S) = \log \frac{\text{Prob}(\text{Data}|H_1(S))}{\text{Prob}(\text{Data}|H_0)}$, corresponds to the generalized likelihood ratio test (GLRT) to verify the null hypothesis (H_0): $\mathbf{c}(v) \sim \mathcal{D}_1, \forall v \in \mathbb{V}$, where \mathcal{D}_1 refers to a predefined background distribution, against the alternative hypothesis ($H_1(S)$): $\mathbf{c}(v) \sim \mathcal{D}_2, \forall v \in S$ and $\mathbf{c}(v) \sim \mathcal{D}_1, \forall v \in \mathbb{V} \setminus S$, where \mathcal{D}_2 refers to a predefined signal distribution. The detection problem is formulated as

$$\min_{S \subseteq \mathbb{V}} -F(S) \quad \text{s.t.} \quad |S| \leq k \text{ and } S \text{ is connected,} \quad (12)$$

where k is a predefined bound on the size of S .

Taking elevated mean scan (EMS) statistic for instance, it aims to decide between $H_0 : \mathbf{c}(v) \sim \mathcal{N}(0, 1), \forall v \in \mathbb{V}$ and $H_1(S) : \mathbf{c}(v) \sim \mathcal{N}(\mu, 1), \forall v \in S$ and $\mathbf{c}(v) \sim \mathcal{N}(0, 1), \forall v \in \mathbb{V} \setminus S$, where for simplicity each node v only has a univariate feature $c(v) \in \mathbb{R}$. This statistic is popularly used for detecting signals among node-level numerical features on graph [Qian et al., 2014; Arias-Castro et al., 2011] and is formulated as $F(S) = (\sum_{v \in S} c(v))^2 / |S|$. Let the vector form of S be $x \in \{0, 1\}^n$, such that $\text{supp}(x) = S$. The connected subgraph detection problem can be reformulated as

$$\min_{x \in \{0, 1\}^n} -(\mathbf{c}^T \mathbf{x})^2 / (\mathbf{1}^T \mathbf{x}) \quad \text{s.t.} \quad \text{supp}(\mathbf{x}) \in \mathbb{M}(k, g = 1), \quad (13)$$

where $\mathbf{c} = [c(1), \dots, c(n)]^T$. To apply GRAPH-MP, we relax the input domain of \mathbf{x} such that $x \in [0, 1]^n$, and the connected subset of nodes can be found as $S = \text{supp}(\mathbf{x}^*)$, the support set of the estimate \mathbf{x}^* that minimizes the strongly convex function [Bach, 2011]:

$$\min_{x \in \mathbb{R}^n} f(\mathbf{x}) = -(\mathbf{c}^T \mathbf{x})^2 / (\mathbf{1}^T \mathbf{x}) + \frac{1}{2} \mathbf{x}^T \mathbf{x} \quad \text{s.t.} \quad \text{supp}(\mathbf{x}) \in \mathbb{M}(k, 1).$$

Assume that \mathbf{c} is normalized, and hence $0 \leq c_i < 1, \forall i$. Let $\hat{c} = \max\{c_i\}$. The Hessian matrix of the above objective function $\nabla^2 f(\mathbf{x}) \succ 0$ and satisfies the inequalities:

$$(1 - \hat{c}^2) \cdot \mathbf{I} \preceq \mathbf{I} - (\mathbf{c} - \frac{\mathbf{c}^T \mathbf{x}}{\mathbf{1}^T \mathbf{x}} \mathbf{1})(\mathbf{c} - \frac{\mathbf{c}^T \mathbf{x}}{\mathbf{1}^T \mathbf{x}} \mathbf{1})^T \preceq \mathbf{I} \cdot \mathbf{I}. \quad (14)$$

According to Lemma 1 (b) in [Yuan et al., 2013]), the objective function $f(\mathbf{x})$ satisfies condition $(\xi, \delta, \mathbb{M}(8k, g))$ -WRSC that $\delta = \sqrt{1 - 2\xi(1 - \hat{c}^2) + \xi^2}$, for any ξ such that $\xi < 2(1 - \hat{c}^2)$. Hence, the geometric convergence of GRAPH-MP as shown in Theorem 4.1 is guaranteed. We note that not all the graph scan statistic functions satisfy the WRSC condition, but, as shown in our experiments, GRAPH-MP works empirically well for all the scan statistic functions tested, and the maximum number of iterations to convergence for optimizing each of these scan statistic functions was less than 10.

We note that our proposed method GRAPH-MP is also applicable to general sparse learning problems (e.g. sparse logistic regression, sparse principle component analysis) subject to graph-structured constraints, and to a variety of subgraph detection problems, such as the detection of anomalous subgraphs, bursty subgraphs, heaviest subgraphs, frequent subgraphs or communication motifs, predictive subgraphs, and compression subgraphs. More details can be found at [Chen, 2016].

6 Experiments

This section evaluates the effectiveness and efficiency of the proposed GRAPH-MP approach for connected subgraph detection. The implementation of GRAPH-MP is available at <https://github.com/baojianzhou/Graph-MP>.

6.1 Experiment Design

Datasets: 1) **Water Pollution Dataset.** The Battle of the Water Sensor Networks (BWSN) [Ostfeld et al., 2008] provides a real-world network of 12,527 nodes and 14831 edges, and 4 nodes with chemical contaminant plumes that are distributed in four different areas. The spreads of contaminant plumes were simulated using the water network simulator EPANET for 8 hours. For each hour, each node has a sensor that reports 1 if it is polluted; 0, otherwise. We randomly selected K percent nodes, and flipped their sensor binary values in order to test the robustness of methods to noises, where $K \in \{2, 4, 6, 8, 10\}$. **The objective is to detect the set of polluted nodes.** 2) **High-Energy Physics Citation Network.** The CitHepPh (high energy physics phenomenology) citation graph is from the e-print arXiv and covers all the citations within a dataset of 34,546 papers with 421,578 edges during the period from January 1993 to April 2002. Each paper is considered as a node, each citation is considered as a edge (direction is not considered), and each node has one attribute denoting the number of citations in a specific year ($t = 1993, \dots, t = 2002$), and another attribute denoting the average number of citations in that year. **The objective is to detect a connected subgraph of nodes (papers) whose citations are abnormally high in comparison with the citations of nodes outside the subgraph.** This subgraph is considered an indicator of a potential emerging research area. The data before 1999 is considered as training data, and the data from 1999 to 2002 is considered as testing data.

Graph Scan Statistics: Three graph scan statistics were considered, including Kulldorff's scan statistic [Neill, 2012], expectation-based Poisson scan statistic (EBP) [Neill, 2012], and elevated mean scan statistic (EMS, Equation (13)) [Qian




Figure 1: Evolving curves of graph scan statistic scores between our method and GenFusedLasso.

	WaterNetwork				CitHepPh			
	Kulldorff	EMS	EBP	Run Time (sec.)	Kulldorff	EMS	EBP	Run Time
Our Method	1668.14	499.97	4032.48	40.98	13859.12	142656.84	9494.62	97.21
GenFusedLasso	541.49	388.04	3051.22	901.51	2861.6	60952.57	6472.84	947.07
EdgeLasso	212.54	308.11	1096.17	70.06	39.42	2.0675.89	261.71	775.61
GraphLaplacian	272.25	182.95	928.41	228.45	1361.91	29463.52	876.31	2637.65
LTSS	686.78	479.40	1733.11	1.33	11965.14	137657.99	9098.22	6.93
EventTree	1304.4	744.45	3677.58	99.27	10651.23	127362.57	8295.47	100.93
AdditiveGraphScan	1127.02	761.08	2794.66	1985.32	12567.29	140078.86	9282.28	2882.74
DepthFirstGraphScan	1059.10	725.65	2674.14	8883.56	7148.46	62774.57	4171.47	9905.45
NPHGS	686.78	479.40	1733.11	1339.46	12021.85	137963.5	9118.96	1244.80

Table 1: Comparison on scores of the three graph scan statistics based on connected subgraphs returned by comparison methods. EMS and EBP refer to Elevated Mean Scan Statistic and Expectation-Based Poisson Statistic, respectively.

et al., 2014]. The first two require that each node has an observed count of events at that node, and an expected count. For the water network dataset, the report of the sensor (0 or 1) at each node is considered as the observed count, and the noise ratio is considered as the expected count. For the CiteHepPh dataset, the number of citations is considered as the observed count, and the average number of citations is considered as the expected count. For the EMS statistic, we consider the ratio of observed and expected counts as the feature.

Comparison Methods: Seven state-of-the-art baseline methods are considered, including EdgeLasso [Sharpnack *et al.*, 2012b], GraphLaplacian [Sharpnack *et al.*, 2012a], LinearTimeSubsetScan (LTSS) [Neill, 2012], EventTree [Rozenshtein *et al.*, 2014], AdditiveGraphScan [Speakman *et al.*, 2013], DepthFirstGraphScan [Speakman *et al.*, 2016], and NPHGS [Chen and Neill, 2014]. We followed strategies recommended by authors in their original papers to tune the related model parameters. Specifically, for EventTree and Graph-Laplacian, we tested the set of λ values: $\{0.02, 0.04, \dots, 2.0\}$. DepthFirstScan is an exact search algorithm and has an exponential time cost in the worst case scenario. We hence set a constraint on the depth of its search to 10 in order to reduce its time complexity.

We also implemented the generalized fused lasso model (GenFusedLasso) for these three graph scan statistics using the framework of alternating direction method of multipliers (ADMM). GenFusedLasso is formalized as

$$\min_{\mathbf{x} \in \mathbb{R}^n} -f(\mathbf{x}) + \lambda \sum_{(i,j) \in \mathbb{E}} \|x_i - x_j\|, \quad (15)$$

where $f(\mathbf{x})$ is a predefined graph scan statistic and the trade-off parameter λ controls the degree of smoothness of neighboring entries in \mathbf{x} . We applied the heuristic rounding step proposed in [Qian *et al.*, 2014] to the numerical vector \mathbf{x} to identify the connected subgraph. We tested the λ values: $\{0.02, 0.04, \dots, 2.0, 5.0, 10.0\}$ and returned the best result.

Our Proposed Method GRAPH-MP: Our proposed GRAPH-MP has a single parameter k , an upper bound of the subgraph size. We set $k = 1000$ by default, as the sizes of subgraphs of interest are often small; otherwise, the detection problem could be less challenging. We note that, to obtain the best performance of our proposed method GRAPH-MP, we should try a set of different k values ($k = 50, 100, 200, 300, \dots, 1000$) and return the best.

Performance Metrics: The overall scores of the three graph scan statistics of the connected subgraphs returned by the competitive methods were compared and analyzed. The objective is to identify methods that could find the connected subgraphs with the largest scores. The running times of different methods are compared.

6.2 Evolving Curves of Graph Scan Statistics

Figure 1 presents the comparison between our method and GenFusedLasso on the scores of the best connected subgraphs that are identified at different iterations based on the Kulldorff’s scan statistic and the EMS statistic. Note that, a heuristic rounding process as proposed in [Qian *et al.*, 2014] was applied to the numerical vector \mathbf{x}^i estimated by GenFusedLasso in order to identify the best connected subgraph at each iteration i . As the setting of the parameter λ

will influence the quality of the detected connected subgraph, the results based on different λ values are also shown in Figure 1. We observe that our proposed method GRAPH-MP converged in less than 5 steps and the qualities (scan statistics scores) of the connected subgraphs identified GRAPH-MP at different iterations were consistently higher than those returned by GenFusedLasso.

6.3 Comparison on Optimization Quality

The comparison between our method and the other eight baseline methods is shown in Table 1. The scores of the three graph scan statistics of the connected subgraphs returned by these methods are reported in this table. The results indicate that our method outperformed all the baseline methods on the scores, except that AdditiveGraphScan achieved the highest EMS score (761.08) on the water network data set. Although AdditiveGraphScan performed reasonably well in overall, this algorithm is a heuristic algorithm and does not have theoretical guarantees.

6.4 Comparison on Time Cost

Table 1 shows the time costs of all competitive methods on the two benchmark data sets. The results indicate that our method was the second fastest algorithm over all the comparison methods. In particular, the running times of our method were 10+ times faster than the majority of the methods.

7 Conclusion and Future Work

This paper presents, GRAPH-MP, an efficient algorithm to minimize a general nonlinear function subject to graph-structured sparsity constraints. For the future work, we plan to explore graph-structured constraints other than connected subgraphs, and analyze theoretical properties of GRAPH-MP for cost functions that do not satisfy the WRSC condition.

A Proof of Theorem 4.1

Proof. Let $\mathbf{r}^{i+1} = \mathbf{x}^{i+1} - \mathbf{x}$. $\|\mathbf{r}^{i+1}\|_2$ is upper bounded as

$$\begin{aligned} \|\mathbf{r}^{i+1}\|_2 = \|\mathbf{x}^{i+1} - \mathbf{x}\|_2 &\leq \|\mathbf{x}^{i+1} - \mathbf{b}\|_2 + \|\mathbf{x} - \mathbf{b}\|_2 \\ &\leq c_T \|\mathbf{x} - \mathbf{b}\|_2 + \|\mathbf{x} - \mathbf{b}\|_2 \\ &\leq (1 + c_T) \|\mathbf{x} - \mathbf{b}\|_2, \end{aligned}$$

which follows from the definition of tail approximation. The component $\|(\mathbf{x} - \mathbf{b})_\Omega\|_2^2$ is upper bounded as

$$\begin{aligned} \|(\mathbf{x} - \mathbf{b})_\Omega\|_2^2 &= \langle \mathbf{b} - \mathbf{x}, (\mathbf{b} - \mathbf{x})_\Omega \rangle \\ &= \langle \mathbf{b} - \mathbf{x} - \xi \nabla_\Omega f(\mathbf{b}) + \xi \nabla_\Omega f(\mathbf{x}), (\mathbf{b} - \mathbf{x})_\Omega \rangle - \\ &\quad \langle \xi \nabla_\Omega f(\mathbf{x}), (\mathbf{b} - \mathbf{x})_\Omega \rangle \\ &\leq \delta \|\mathbf{b} - \mathbf{x}\|_2 \|(\mathbf{b} - \mathbf{x})_\Omega\|_2 + \xi \|\nabla_\Omega f(\mathbf{x})\|_2 \|(\mathbf{b} - \mathbf{x})_\Omega\|_2, \end{aligned}$$

where the second equality follows from the fact that $\nabla_\Omega f(\mathbf{b}) = 0$ since \mathbf{b} is the solution to the problem in Step 6 of Algorithm 1, and the last inequality follows from condition $(\xi, \delta, \mathbb{M}(8k, g))$ -WRSC. After simplification, we have

$$\|(\mathbf{x} - \mathbf{b})_\Omega\|_2 \leq \delta \|\mathbf{b} - \mathbf{x}\|_2 + \xi \|\nabla_\Omega f(\mathbf{x})\|_2.$$

It follows that

$$\begin{aligned} \|(\mathbf{x} - \mathbf{b})\|_2 &\leq \|(\mathbf{x} - \mathbf{b})_\Omega\|_2 + \|(\mathbf{x} - \mathbf{b})_{\Omega^c}\|_2 \\ &\leq \delta \|\mathbf{b} - \mathbf{x}\|_2 + \xi \|\nabla_\Omega f(\mathbf{x})\|_2 + \|(\mathbf{x} - \mathbf{b})_{\Omega^c}\|_2. \end{aligned}$$

After rearrangement we obtain

$$\begin{aligned} \|\mathbf{b} - \mathbf{x}\|_2 &\leq \frac{\|(\mathbf{b} - \mathbf{x})_{\Omega^c}\|_2}{1 - \delta} + \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta} \\ &= \frac{\|\mathbf{x}_{\Omega^c}\|_2}{1 - \delta} + \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta} = \frac{\|(\mathbf{x} - \mathbf{x}^i)_{\Omega^c}\|_2}{1 - \delta} + \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta} \\ &= \frac{\|\mathbf{r}_{\Omega^c}\|_2}{1 - \delta} + \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta} \leq \frac{\|\mathbf{r}_{\Gamma^c}^i\|_2}{1 - \delta} + \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta}, \end{aligned}$$

where the first equality follows from the fact that $\text{supp}(\mathbf{b}) \subseteq \Omega$, the second and last inequalities follow from the fact that $\Omega = \Gamma \cup \text{supp}(\mathbf{x}^i)$. Combining above inequalities, we obtain

$$\|\mathbf{r}^{i+1}\|_2 \leq (1 + c_T) \frac{\|\mathbf{r}_{\Omega^c}^i\|_2}{1 - \delta} + (1 + c_T) \frac{\xi \|\nabla_\Omega f(\mathbf{x})\|_2}{1 - \delta}.$$

From Lemma A.2, we have

$$\|\mathbf{r}_{\Gamma^c}^i\|_2 \leq \sqrt{1 - \eta^2} \|\mathbf{r}^i\|_2 + \left[\frac{\xi(1 + c_H)}{\eta} + \frac{\xi\eta(1 + c_H)}{\sqrt{1 - \eta^2}} \right] \|\nabla_I f(\mathbf{x})\|_2$$

Combining the above inequalities, we prove the theorem. \square

Lemma A.1. [Yuan et al., 2013] Assume that f is a differentiable function. If f satisfies condition $(\xi, \delta, \mathbb{M})$ -WRSC, then $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\text{supp}(\mathbf{x}) \cup \text{supp}(\mathbf{y}) \subset S \in \mathbb{M}$, the following two inequalities hold

$$\begin{aligned} \frac{1 - \delta}{\xi} \|\mathbf{x} - \mathbf{y}\|_2 &\leq \|\nabla_S f(\mathbf{x}) - \nabla_S f(\mathbf{y})\|_2 \leq \frac{1 + \delta}{\xi} \|\mathbf{x} - \mathbf{y}\|_2, \\ f(\mathbf{x}) &\leq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \frac{1 + \delta}{2\xi} \|\mathbf{x} - \mathbf{y}\|_2^2. \end{aligned}$$

Lemma A.2. Let $\mathbf{r}^i = \mathbf{x}^i - \mathbf{x}$ and $\Gamma = H(\nabla f(\mathbf{x}^i))$. Then

$$\|\mathbf{r}_{\Gamma^c}^i\|_2 \leq \sqrt{1 - \eta^2} \|\mathbf{r}^i\|_2 + \left[\frac{\xi(1 + c_H)}{\eta} + \frac{\xi\eta(1 + c_H)}{\sqrt{1 - \eta^2}} \right] \|\nabla_I f(\mathbf{x})\|_2,$$

where $\eta = c_H(1 - \delta) - \delta$ and $I = \arg \max_{S \in \mathbb{M}(8k, g)} \|\nabla_S f(\mathbf{x})\|_2$. We assume that c_H and δ are such that $\eta > 0$.

Proof. Denote $\Phi = \text{supp}(\mathbf{x}) \in \mathbb{M}(k, g)$, $\Gamma = H(\nabla f(\mathbf{x}^i)) \in \mathbb{M}(2k, g)$, $\mathbf{r}^i = \mathbf{x}^i - \mathbf{x}$, and $\Omega = \text{supp}(\mathbf{r}^i) \in \mathbb{M}(6k, g)$. The component $\|\nabla_\Gamma f(\mathbf{x}^i)\|_2$ can be lower bounded as

$$\begin{aligned} \|\nabla_\Gamma f(\mathbf{x}^i)\|_2 &\geq c_H (\|\nabla_\Phi f(\mathbf{x}^i) - \nabla_\Phi f(\mathbf{x})\|_2 - \|\nabla_\Phi f(\mathbf{x})\|_2) \\ &\geq c_H(1 - \delta) \|\mathbf{r}^i\|_2 / \xi - c_H \|\nabla_I f(\mathbf{x})\|_2, \end{aligned}$$

where the last inequality follows from Lemma A.1. The component $\|\nabla_\Gamma f(\mathbf{x}^i)\|_2$ can also be upper bounded as

$$\begin{aligned} \|\nabla_\Gamma f(\mathbf{x}^i)\|_2 &\leq \frac{1}{\xi} \|\xi \nabla_\Gamma f(\mathbf{x}^i) - \xi \nabla_\Gamma f(\mathbf{x})\|_2 + \|\nabla_\Gamma f(\mathbf{x})\|_2 \\ &\leq \frac{1}{\xi} \|\xi \nabla_\Gamma f(\mathbf{x}^i) - \xi \nabla_\Gamma f(\mathbf{x}) - \mathbf{r}_\Gamma^i + \mathbf{r}_\Gamma^i\|_2 + \|\nabla_\Gamma f(\mathbf{x})\|_2 \\ &\leq \frac{1}{\xi} \|\xi \nabla_{\Gamma \cup \Omega} f(\mathbf{x}^i) - \xi \nabla_{\Gamma \cup \Omega} f(\mathbf{x}) - \mathbf{r}_{\Gamma \cup \Omega}^i\|_2 + \|\mathbf{r}_\Gamma^i\|_2 + \|\nabla_\Gamma f(\mathbf{x})\|_2 \\ &\leq \frac{\delta}{\xi} \cdot \|\mathbf{r}^i\|_2 + \frac{1}{\xi} \|\mathbf{r}_\Gamma^i\|_2 + \|\nabla_I f(\mathbf{x})\|_2, \end{aligned}$$

where the last inequality follows from condition $(\xi, \delta, \mathbb{M}(8k, g))$ -WRSC and the fact that $\mathbf{r}_{\Gamma \cup \Omega}^i = \mathbf{r}^i$. Let $\eta = (c_H \cdot (1 - \delta) - \delta)$. Combining the two bounds and grouping terms, we have $\|\mathbf{r}_\Gamma^i\|_2 \geq \eta \|\mathbf{r}^i\|_2 - \xi(1 + c_H) \|\nabla_I f(\mathbf{x})\|_2$. After a number of algebraic manipulations similar to those used in [Hegde et al., 2014a] Page 11, we prove the lemma. Readers are referred to [Chen, 2016] for more details of the proof. \square

B Acknowledgements

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract D12PC00337. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the US government.

References

- [Arias-Castro *et al.*, 2011] Ery Arias-Castro, Emmanuel J Candès, and Arnaud Durand. Detection of an anomalous cluster in a network. *ANN STAT*, pp 278–304, 2011.
- [Asteris *et al.*, 2015] Megasthenis Asteris, Anastasios Kyrillidis, Alexandros G Dimakis, Han-Gyol Yi, et al. Stay on path: PCA along graph paths. In *ICML*, 2015.
- [Bach *et al.*, 2012] Francis Bach, Rodolphe Jenatton, et al. Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468, 2012.
- [Bach, 2011] Francis Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends in Mach. Learn.*, 6(2-3), pp 145–373, 2011.
- [Bahmani *et al.*, 2013] Sohail Bahmani, Bhiksha Raj, and Petros T Boufounos. Greedy sparsity-constrained optimization. *JMLR*, 14(1):807–841, 2013.
- [Bahmani *et al.*, 2016] Sohail Bahmani, Petros T Boufounos, and Bhiksha Raj. Learning model-based sparsity via projected gradient descent. *IEEE Trans. Info. Theory*, 62(4):20922099, 2016.
- [Blumensath, 2013] Thomas Blumensath. Compressed sensing with nonlinear observations and related nonlinear optimization problems. *Information Theory, IEEE Transactions on*, 59(6):3466–3474, 2013.
- [Chen and Neill, 2014] Feng Chen and Daniel B. Neill. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *KDD*, pp 1166–1175, 2014.
- [Duarte and Eldar, 2011] Marco F Duarte and Yonina C Eldar. Structured compressed sensing: From theory to applications. *ITSP*, 59(9):4053–4085, 2011.
- [Hegde *et al.*, 2014a] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. Approximation-tolerant model-based compressive sensing. In *SODA*, pp 1544–1561, 2014.
- [Hegde *et al.*, 2014b] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. A fast approximation algorithm for tree-sparse recovery. In *ISIT*, pp 1842–1846, 2014.
- [Hegde *et al.*, 2015a] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. Fast algorithms for structured sparsity. *Bulletin of EATCS*, 3(117), 2015.
- [Hegde *et al.*, 2015b] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt. A nearly-linear time framework for graph-structured sparsity. In *ICML*, pp 928–937, 2015.
- [Huang *et al.*, 2011] Junzhou Huang, Tong Zhang, and Dimitris Metaxas. Learning with structured sparsity. *The Journal of Machine Learning Research*, 12:3371–3412, 2011.
- [Jacob *et al.*, 2009] Laurent Jacob, Guillaume Obozinski, and Jean-Philippe Vert. Group lasso with overlap and graph lasso. In *ICML*, pp 433–440, 2009.
- [Jain *et al.*, 2014] Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional m-estimation. In *NIPS*, 2014.
- [Neill, 2012] Daniel B Neill. Fast subset scan for spatial pattern detection. *JRSS: Series B*, 74(2):337–360, 2012.
- [Ostfeld *et al.*, 2008] Avi Ostfeld, James G Uber, et al. The battle of the water sensor networks (BWSN): A design challenge for engineers and algorithms. *JWRPM*, 134(6):556–568, 2008.
- [Qian *et al.*, 2014] J. Qian, V. Saligrama, and Y. Chen. Connected sub-graph detection. In *AISTATS*, 2014.
- [Rozenstein *et al.*, 2014] Polina Rozenstein, Aris Anagnostopoulos, Aristides Gionis, and Nikolaj Tatti. Event detection in activity networks. In *KDD*, 2014.
- [Sharpnack *et al.*, 2012a] James Sharpnack, Alessandro Rinaldo, and Aarti Singh. Change-point detection over graphs with the spectral scan statistic. *arXiv preprint arXiv:1206.0773*, 2012.
- [Sharpnack *et al.*, 2012b] James Sharpnack, Alessandro Rinaldo, and Aarti Singh. Sparsistency of the edge lasso over graphs. 2012.
- [Speakman *et al.*, 2016] Skyler Speakman, Edward McFowland Iii, and Daniel B. Neill. Scalable detection of anomalous patterns with connectivity constraints. *Journal of Computational and Graphical Statistics*, 2016.
- [Speakman *et al.*, 2013] S. Speakman, Y. Zhang, and D. B. Neill. Dynamic pattern detection with temporal consistency and connectivity constraints. In *ICDM*, 2013.
- [Tewari *et al.*, 2011] A. Tewari, P. K Ravikumar, and I. S Dhillon. Greedy algorithms for structurally constrained high dimensional problems. In *NIPS*, pp 882–890, 2011.
- [Xin *et al.*, 2014] Bo Xin, Yoshinobu Kawahara, Yizhou Wang, and Wen Gao. Efficient generalized fused lasso and its application to the diagnosis of alzheimers disease. In *AAAI*, pp 2163–2169, 2014.
- [Yuan and Liu, 2014] Xiao-Tong Yuan and Qingshan Liu. Newton greedy pursuit: A quadratic approximation method for sparsity-constrained optimization. In *CVPR*, pp 4122–4129, 2014.
- [Yuan *et al.*, 2013] Xiao-Tong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit for sparsity-constrained optimization. In *ICML*, 2014.
- [Zhang, 2009] Tong Zhang. Adaptive forward-backward greedy algorithm for sparse learning with linear models. In *NIPS*, pp 1921–1928, 2009.
- [Chen, 2016] Feng Chen and Baojian Zhou. A generalized matching pursuit approach for graph-structured sparsity (extended version). In *arXiv*, 2016.