

---

# Detecting Weak but Hierarchically-Structured Patterns in Networks

---

**Aarti Singh**

Machine Learning Department  
Carnegie Mellon University  
aartisinhg@cmu.edu

**Robert Nowak**

Electrical and Computer Engineering  
University of Wisconsin - Madison  
nowak@engr.wisc.edu

**Robert Calderbank**

Electrical Engineering, PACM  
Princeton University  
calderbk@princeton.edu

## Abstract

The ability to detect weak distributed activation patterns in networks is critical to several applications, such as identifying the onset of anomalous activity or incipient congestion in the Internet, or faint traces of a biochemical spread by a sensor network. This is a challenging problem since weak distributed patterns can be invisible in per node statistics as well as a global network-wide aggregate. Most prior work considers situations in which the activation/non-activation of each node is statistically independent, but this is unrealistic in many problems. In this paper, we consider structured patterns arising from statistical dependencies in the activation process. Our contributions are three-fold. First, we propose a sparsifying transform that succinctly represents structured activation patterns that conform to a hierarchical dependency graph. Second, we establish that the proposed transform facilitates detection of very weak activation patterns that cannot be detected with existing methods. Third, we show that the structure of the hierarchical dependency graph governing the activation process, and hence the network transform, can be learnt from very few (logarithmic in network size) independent snapshots of network activity.

## 1 Introduction

Consider the problem of detecting a weak binary pattern corrupted by noise that is observed at  $p$  network nodes:

$$y_i = \mu x_i + \epsilon_i \quad i = 1, \dots, p$$

Here  $y_i$  denotes the observation at node  $i$  and  $\mathbf{x} = [x_1, \dots, x_p] \in \{0, 1\}^p$  is the  $p$ -dimensional *unknown* binary activation pattern,  $\mu > 0$  denotes an *unknown* signal strength, and the noise  $\epsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ , the Gaussian distribution with mean zero and variance  $\sigma^2$ . The condition

---

Appearing in Proceedings of the 13<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2010, Chia Laguna Resort, Sardinia, Italy. Volume 9 of JMLR: W&CP 9. Copyright 2010 by the authors.

$x_i = 0, i = 1, \dots, p$ , is the baseline or normal operating condition (no signal present). If  $x_i > 0$  for one or more  $i$ , then a signal or activation is present in the network. We are interested not in arbitrary patterns of activation, but rather our focus is on patterns that are related to the physical structure of the network and/or to other statistical dependencies in the signal. This is motivated by problems arising in practice, as discussed below. Specifically, we consider classes of patterns that are supported over hierarchically-structured groups or clusters of nodes. Such a hierarchical structure could arise due to the physical topology of the network and/or due to dependencies between the nodes. For example, hierarchical dependencies are known to exist in gene networks due to shared regulatory pathways (Yu & Gerstein 2006, Girvan & Newman 2002), empirical studies show that Internet path properties such as delay and bandwidth are well-approximated by tree-embeddings (Ramasubramanian, Malkhi, Kuhn, Balakrishnan & Akella 2009), sensor networks are often hierarchically structured for efficient management (Sankaranarayanan, Kramer & Mandayam 2004), and communities in social networks can be hierarchical (Girvan & Newman 2002). We address the problem of detecting the presence of weak but hierarchically-structured activation patterns in the network. This problem is of interest in several applications including detecting incipient congestion or faint traces of malicious activity in the Internet, early detection of a chemical spread or bio-hazard by a sensor network, identification of differentially expressed genes in microarray data, or malicious groups in social networks.

If  $\mathbf{x}$  is known, then the optimal detector is based on aggregating the measurements of the locations known to contain the signal (e.g., in the classical distributed detection literature it is often assumed that  $x_i = 1$  for all  $i$  or  $x_i = 0$  for all  $i$  (Varshney 1996)). We are interested in cases where  $\mathbf{x}$  is unknown. If  $\mathbf{x}$  is arbitrary, this is a problem in literature known as the *multi-channel signal detection problem* (Ingster & Suslina 2002). In this case, global aggregation rule (testing the average of all node measurements) can reliably detect any signal strength  $\mu > 0$  if the number of active locations  $\|\mathbf{x}\|_0 > \sqrt{p}$ . This is because  $\frac{1}{\sqrt{p}} \sum_{i=1}^p y_i \sim \mathcal{N}\left(\mu \frac{\|\mathbf{x}\|_0}{\sqrt{p}}, \sigma^2\right)$ , and therefore as the net-

work size  $p$  grows, the probability of false alarm and miss can be driven to zero by choosing an appropriate threshold. However, in the high-dimensional setting when  $p$  is very large and the activation is sparse  $\|\mathbf{x}\|_0 \leq \sqrt{p}$ , different approaches to detection are required. If the signal strength  $\mu > \sqrt{2\sigma^2 \log p}$ , then the signal can be reliably detected using the max statistic  $\max_i y_i$ , irrespective of the signal sparsity level. This is because if there is no signal, the max statistic due to noise alone (maximum of  $p$  iid  $\mathcal{N}(0, \sigma^2)$  random variables) is  $\leq \sqrt{2\sigma^2 \log p}$  with probability 1, in the large  $p$  limit. Therefore, the most challenging case is when the network activation is

$$\text{weak: } \mu < \sqrt{2\sigma^2 \log p} \quad \text{and} \quad \text{sparse: } \|\mathbf{x}\|_0 < \sqrt{p}$$

In this case, the signal is buried in noise and cannot be detected in per node measurement or in global network-wide aggregate. This necessitates *selective and adaptive fusion* where the node measurements to be aggregated are chosen in a data-driven fashion. One approach that is common in the signal processing literature is to consider the generalized likelihood ratio test (GLRT) statistic,  $\max_{\mathbf{x} \in \{0,1\}^p} \mathbf{x}^T \mathbf{y} / \mathbf{x}^T \mathbf{x}$ , where the observed vector is matched with all  $2^p$  possible true activation patterns. However, in high-dimensional settings, the GLRT is computationally intractable. For weak and sparse signals, the limits of detectability were studied by Ingster (Ingster & Suslina 2002), and subtle tests that are adaptive in various ranges of the unknown sparsity level were investigated. More recently, test statistics have been proposed (Jin & Donoho 2004, Jager & Wellner 2007) that can attain the detection boundary simultaneously for any unknown sparsity level. Also see (Ingster, Pouet & Tsybakov 2009) for a generalization of this problem. However, all this work assumes that activations at nodes are independent of each other. As a result, the signal strength  $\mu$  must be  $> c\sqrt{\log p}$  for a constant  $c > 0$  and the signal cannot be too weak.

The assumption of independent activations is often unreasonable in a network setting, where the observations at nodes tend to be highly dependent due to the structure of the network and/or dependencies in the activation process itself. For example, routers in the same autonomous system will show similar variations in round-trip-time measurements, or co-located sensors monitoring an environmental phenomena will have correlated measurements. Recently, there has been some work aimed at structured patterns of activation in graphs (A.-Castro, Donoho & Huo 2005, A.-Castro, Candés, Helgason & Zeitouni 2007, A.-Castro, Candés & Durand 2010), which indicates that it is possible to detect even weaker signals by leveraging the statistical dependencies in the activation process. Of these, the lattice-based models in (A.-Castro et al. 2010) are most closely related to our work, but they do not capture the hierarchical structure we have in mind, nor do they appear to offer a computationally tractable approach to detection. We also mention the recent work of (A.-Berry, Broutin,

Devroye & Lugosi 2009), which establishes fundamental limits of detectability for several classes of structured patterns in graphs. The detection tests proposed in that paper are generally combinatorial in nature (like the GLRT mentioned above) requiring a brute-force examination of all patterns in each class, and therefore are computationally prohibitive except in very low-dimensional cases.

In this paper, we consider a different class of patterns that reflects the hierarchical dependencies present in many real-world networks and leads to computationally practical detection methods. Furthermore, we demonstrate that it is possible to learn the hierarchical dependency structure of the class from a relatively small number of observations, adding to the practical potential of our framework. The hierarchical dependencies structures we consider tend to result in network activation patterns that are supported over hierarchically-organized groups or clusters of nodes. We will show that such *structured* activation patterns can be sparsified even further by an orthonormal transformation that is adapted to the dependency structure. The transform concentrates the unknown  $\mathbf{x}$  in a few large basis coefficients, thus facilitating detection. We show that if the canonical domain sparsity  $\|\mathbf{x}\|_0 \sim p^{1-\alpha}$  and the transform domain sparsity scales as  $p^{1-\beta}$ , where  $\beta > \alpha$ , then the threshold of detection scales as  $\mu > \sqrt{2p^{-(\beta-\alpha)}\sigma^2 \log p}$ . Contrasting this with the detectability threshold of earlier methods  $\mu > \sqrt{2\eta_\alpha\sigma^2 \log p}$  (Ingster & Suslina 2002, Jin & Donoho 2004) (where  $0 < \eta_\alpha < 1$  is independent of  $p$ ), we see that a polynomial improvement is attained if the activation pattern is sparser in transform domain. By exploiting the structure of  $\mathbf{x}$ , we can detect extremely faint activations that can not be detected with existing methods.

Our contributions are three-fold. First, we propose a sparsifying transform based on hierarchical clustering that is adapted to the dependency structure of network measurements. We propose a practically-motivated generative model that allows for arbitrary activation patterns, but favors patterns that are supported over hierarchically-organized groups of nodes. We show that patterns from this model are compressed by the sparsifying transform. Though we focus on the detection problem in this paper, the sparsifying transform could be exploited in other problem domains, e.g. de-noising, compression, sparse regression, variable selection, etc. Second, we establish that the sparsifying transform can amplify very weak activation patterns by effectively performing adaptive fusion of the network measurements. Since the network activity is summarized in a few large transform coefficients, the signal-to-noise ratio (SNR) is increased, and this facilitates detection of very weak activation patterns. We quantify the improvement in the detection threshold relative to existing methods. The detection method we propose is a constructive procedure and computationally efficient. Third, we do not necessarily assume that the graph structure is known a priori, and show

that the dependency structure, and hence the sparsifying transform, can be learnt from very few,  $O(\log p)$ , multiple independent snapshots of network measurements.

We introduce the sparsifying transform in section 2, a generative model for hierarchically-structured patterns and corresponding sparsifying properties and detection threshold in section 3, sample complexity of learning the transform in Section 4 and simulations in Section 5. Proofs sketches are given in the Appendix.

## 2 Hierarchical structure in Networks

As discussed in the introduction, activation patterns in large-scale networks such as the Internet, sensor, biological and social networks often have hierarchical dependencies. This hierarchical dependency structure can be exploited to enable detection of very weak and sparse patterns of activity. In this section, we propose a transform that is adapted to a given set of pairwise similarities between nodes. The similarity of node  $i$  and  $j$  is denoted by  $r_{ij}$ . For example,  $r_{ij}$  could be the covariance between measurements at nodes  $i$  and  $j$ , but other similarity measures can also be employed. The transform is derived from a hierarchical clustering based on the similarity matrix  $\{r_{ij}\}$ . If the matrix reflects an underlying hierarchical dependency structure, then the resulting transform sparsifies activation patterns supported on hierarchically-organized groups of nodes.

### 2.1 Hierarchical Clustering of Nodes

We employ a standard, bottom-up agglomerative clustering algorithm. The algorithm takes as input a set of pairwise similarities  $\{r_{ij}\}$  and returns a hierarchical set of clusters/groups of nodes, denoted as  $\mathcal{H}$ . The algorithm is described in Figure 1. Suppose instead that we are given a hierarchical set of clusters  $\mathcal{H}^*$ . What conditions must a similarity matrix satisfy, in relation to  $\mathcal{H}^*$ , so that the agglomerative clustering algorithm recovers  $\mathcal{H}^*$  and not some other hierarchical clusters? This is an important question for several reasons as we will see in subsequent sections (e.g., to robustly identify  $\mathcal{H}^*$  from a noisy observation of the similarity matrix). To answer this question, first note that the agglomerative clustering algorithm always merges two clusters at each step. Therefore, the most we can hope to say is that under some conditions on the similarity matrix, the agglomerative clustering algorithm produces a hierarchical set of clusters  $\mathcal{H}$ , such that  $\mathcal{H}^* \subset \mathcal{H}$ ; i.e.,  $\mathcal{H}$  contains all cluster sets in  $\mathcal{H}^*$ , but may include additional subsets due to the restriction of binary merging. The following lemma gives a sufficient condition on the similarity matrix to guarantee that this is the case. The proof is straightforward and omitted to save space.

**Lemma 1.** *Suppose we are given a collection of hierarchical clusters  $\mathcal{H}^*$ . If for every pair of clusters  $(c, c') \in \mathcal{H}^*$ , where  $c' \subset c$ , the maximum similarity between any  $i \in c'$*

*and  $j \in c/c'$  is smaller than the minimum similarity between any pair of nodes in  $c'$ , then the agglomerative clustering algorithm of Figure 1 recovers  $\mathcal{H}^*$ .*

### 2.2 Hierarchical Basis for Network Patterns

Based on a hierarchical clustering of network nodes, we propose the following unbalanced Haar basis representation for activation patterns. When two clusters  $c_1$  and  $c_2$  are merged in the agglomerative clustering algorithm, a normalized basis vector is defined (up to normalization) by  $\mathbf{b} \propto \frac{1}{|c_2|} \mathbf{1}_{c_2} - \frac{1}{|c_1|} \mathbf{1}_{c_1}$ , where  $\mathbf{1}_{c_i}$  denotes the indicator of the support of subcluster  $c_i$ . Projecting the activation pattern  $\mathbf{x}$  onto this basis vector computes a difference of the average measurement on each constituent cluster. As a result, the basis coefficient  $\mathbf{b}^T \mathbf{x}$  is zero if the nodes in the constituent clusters are all active or inactive. Thus, the basis vectors possess one vanishing moment akin to standard Haar wavelet transform, and will sparsify activation patterns that are constant over the merged clusters. This procedure yields  $p - 1$  difference basis vectors. These basis vectors are augmented with the constant vector that computes the global average. The resulting vectors form the columns of an orthonormal unbalanced Haar transform matrix  $\mathbb{B}$ .

*Input:* Set of all nodes  $\mathcal{L} = \{1, \dots, p\}$  and pairwise similarities  $\{r_{ij}\}_{i,j \in \mathcal{L}}$   
*Initialize:* Clusters  $\mathcal{C} = \{\{1\}, \{2\}, \dots, \{p\}\}$ , Hierarchical clustering  $\mathcal{H} = \mathcal{C}$ , Basis  $\mathbb{B} = []$   
**while**  $|\mathcal{C}| > 1$   
    Select  $(c_1, c_2) = \arg \max_{c_1, c_2 \in \mathcal{C}} \frac{\sum_{i \in c_1} \sum_{j \in c_2} r_{ij}}{|c_1||c_2|}$   
    Merge  $c = c_1 \cup c_2$   
    Update  $\mathcal{H} = \mathcal{H} \cup \{c\}$ ,  $\mathcal{C} = (\mathcal{C}/\{c_1, c_2\}) \cup \{c\}$   
    Construct unbalanced Haar basis vector:  
     $\mathbf{b} = \frac{\sqrt{|c_1||c_2|}}{\sqrt{|c_1| + |c_2|}} \left[ \frac{1}{|c_2|} \mathbf{1}_{c_2} - \frac{1}{|c_1|} \mathbf{1}_{c_1} \right]$ ,  $\mathbb{B} = [\mathbb{B} \mid \mathbf{b}]$   
**end**  
 $\mathbf{b} = \frac{1}{\sqrt{|\mathcal{L}|}} \mathbf{1}_{\mathcal{L}}$ ,  $\mathbb{B} = [\mathbb{B} \mid \mathbf{b}]$   
*Output:*  $\mathbb{B}, \mathcal{H}$

Figure 1: Algorithm for hierarchical clustering.

The proposed method of hierarchical clustering followed by basis construction is similar in spirit to the recent work of Lee et al. (Lee, Nadler & Wasserman 2008) on treelets and of Murtagh (Murtagh 2007). However, treelets do not lead to a sparsifying transform in general if the node measurements or aggregates have different variances. The work of Murtagh uses balanced Haar wavelets on a dendrogram and does not yield an orthonormal basis since the basis vectors are not constant on sub-groups of nodes. As a result, the transform coefficients are correlated and dependent, making the resulting statistics difficult to analyze.

Our procedure, on the other hand, is based on *unbalanced* Haar wavelets which are constant on sub-groups of nodes and thus result in orthogonal vectors.

### 2.3 Activations of Hierarchically-Organized Groups

To illustrate the effectiveness of the proposed transform, consider activation patterns generated by the union of a small number of clusters of the hierarchical collection  $\mathcal{H}$ , i.e. let  $\mathbf{x} = \mathbf{1}_{\cup_{i=1}^m c_i}$ , where  $c_i \in \mathcal{H}$ . Then it is not difficult to see that the transform of  $\mathbf{x}$  will produce no more than  $O(m)$  non-zero basis coefficients. The magnitude of each coefficient will be proportional to the square-root of the number of nodes in the corresponding cluster on which the basis is supported. Suppose that the largest cluster contains  $k$  nodes. Then the largest coefficient of  $\mathbf{x}$  will be on the order of  $\sqrt{k}$ . This implies that the corresponding coefficient of the noisy observations  $\mathbf{y}$  will have a signal-to-noise energy ratio (SNR) of order  $k\mu^2/\sigma^2$ , compared to the per node SNR of  $\mu^2/\sigma^2$  in the canonical domain, making the activation much more easily detectable.

In practice, actual activation patterns may only approximate this sort of ideal condition, but the transform can still significantly boost the SNR even when the underlying activation is only approximately sparse in the transform domain. In the next section we propose a practically-motivated generative model capable of generating arbitrary patterns. As the parameter of the model is varied, the patterns generated from the model tend to have varying degrees of sparseness in the transform domain.

## 3 Sparsifying and Detecting Activations

In this section, we study the sparsifying capabilities of the proposed transform, and the corresponding improvements that can be attained in the detection threshold. For this, we introduce a generative model that, with high probability, produces patterns that are approximately sparse.

### 3.1 A Generative Model for Activations

We model the hierarchical dependencies governing the activation process by a multi-scale latent Ising model, defined as follows. Let  $\mathcal{T}^* = (V, E)$  denote a tree-structured graph with  $V$  as the vertex set and  $E$  as the edge set. For simplicity, we assume that the degree of each node is uniform, denoted as  $d$ , and let  $L = \log_d p$  denote the depth of the tree. The leaves  $\mathcal{L}$  of the tree are at the deepest level  $L$  and correspond to the network nodes, while the internal vertices characterize the multi-scale dependencies between the node measurements. Let  $\mathbf{z}$  denote a  $|V|$ -dimensional vector of variables defined over the complete tree, but we only observe  $\mathbf{x} = \{z_i\}_{i \in \mathcal{L}}$ , the  $p$ -dimensional vector of network observations. We assume that  $\mathbf{z}$  (and hence  $\mathbf{x}$ ) is generated

according to the following probabilistic Ising model:

$$p(\mathbf{z}) \propto \exp \left( \sum_{\ell=1}^L \gamma_\ell \sum_{i \in V_\ell} [z_i z_{\pi(i)} + (1 - z_i)(1 - z_{\pi(i)})] \right)$$

Here  $V_\ell$  denotes the vertices at level  $\ell$ , and  $\gamma_\ell > 0$  characterizes the strength of pairwise interaction between a vertex  $i$  at level  $\ell$  and its parent  $\pi(i)$ . This model implies that the  $2^p$  possible activation patterns are not equiprobable, and the probability of a pattern is higher if the variables agree with their parents in the tree dependency graph  $\mathcal{T}^*$ . This is a natural model for several application domains where the activation is governed by a contact process, e.g. the spread of an infection or disease.

### 3.2 Canonical and Transform Domain Sparsity

To evaluate the transform domain sparsity, we first establish that the latent tree dependency graph  $\mathcal{T}^*$  can be recovered by the agglomerative hierarchical clustering algorithm of Figure 1. Based on a result by Falk (Falk 1975), the covariance between any two leaf variables  $i$  and  $j$  is proportional to  $\prod_{\ell=\ell'+1}^L (\tanh \gamma_\ell)^2$ , where  $\ell'$  denotes the level of the root of the smallest subtree containing  $i$  and  $j$  (i.e. smallest cluster containing  $i$  and  $j$ ). Thus, if the covariance is used as the similarity measure, it is easy to verify that it satisfies the conditions of Lemma 1. This is important since covariance can be estimated from observations of the network.

**Proposition 1.** *The hierarchical clustering algorithm of Figure 1 perfectly recovers the tree-structured dependency graph  $\mathcal{T}^*$  on which the Ising model is defined, using covariance between the leaf variables as the similarity measure.*

We now show how the unbalanced Haar basis built on the tree dependency graph  $\mathcal{T}^*$  leads to a sparse representation of binary patterns drawn from the multi-scale Ising model. Recall that a transform coefficient is zero if the activation pattern is constant over the basis vector's support.

**Theorem 1.** *Consider a pattern  $\mathbf{x}$  drawn at random from a latent Ising model on a tree-structured graph with uniform degree  $d$  and depth  $L = \log_d p$ . If the interaction strength scales with the level  $\ell$  as  $\gamma_\ell = \ell\beta \log d$  where  $0 \leq \beta \leq 1$ , then with probability  $> 1 - \delta$  and for  $p$  large enough, the number of non-zero transform coefficients are bounded by*

$$\|\mathbb{B}^T \mathbf{x}\|_0 \leq 3d(\log_d p)^2 p^{1-\beta}.$$

Proof is given in the Appendix. Since the interaction strength increases with level, variables at deeper levels are less likely to disagree with their parents and activation patterns supported over groups of nodes are favored. The theorem states that, with high probability, patterns generated by this model are approximately sparse in the proposed transform domain. The degree of sparsity is governed by  $\beta$ , the rate at which the interaction strength increases with level.

We also have in mind situations in which the number of total activations in the network is small, i.e.,  $\|\mathbf{x}\|_0 < \sqrt{p}$ , which renders the naive global fusion test statistic unreliable (see discussion in Introduction). To make widespread activations less probable, we constrain the Ising model as follows. Set the root vertex to the value 0. Let  $\ell_0 = \frac{\alpha}{\beta}L$ , where  $0 < \alpha < \beta$ . Let  $\gamma_\ell = \ell\beta \log d$  for  $\ell \geq \ell_0$ , and  $\gamma_\ell = \infty$  for  $\ell < \ell_0$ . This model forces variables at scales coarser than  $\ell_0$  to be identically 0. Proof of the following theorem is given in the Appendix.

**Theorem 2.** *Consider a pattern  $\mathbf{x}$  drawn at random from a latent Ising model on a tree-structured graph with uniform degree  $d$  and depth  $L = \log_d p$ . Let  $\ell_0 = \frac{\alpha}{\beta}L$ , where  $0 < \alpha < \beta$ , and the interaction strength scale with the level  $\ell$  as  $\gamma_\ell = \ell\beta \log d$  for  $\ell \geq \ell_0$ , and  $\gamma_\ell = \infty$  for  $\ell < \ell_0$ . If the pattern corresponds to the root variable taking value zero, then with probability  $> 1 - 4\delta$  and for  $p$  large enough, the number of non-zero transform coefficients are bounded by*

$$\|\mathbb{B}^T \mathbf{x}\|_0 \leq 3d(\log_d p)^2 p^{1-\beta},$$

and the canonical domain sparsity is bounded as

$$cp^{1-\alpha} \leq \|\mathbf{x}\|_0 \leq C(\log_d p)p^{1-\alpha},$$

where  $C > c > 0$  are constant.

The result of the theorem states that the transform domain sparsity scales as  $p^{1-\beta}$  (and is therefore determined by the rate at which the interaction strength increases with level), while the canonical domain sparsity scales as  $p^{1-\alpha}$  (and is therefore determined by the smallest interaction strength between a variable and its parent). Since  $\beta > \alpha$ , the proposed transform enhances the sparsity of canonically sparse patterns that have a hierarchical structure. In the next section, we show that this enhanced sparsity implies a higher Signal-to-Noise (SNR) ratio in the transform domain, thus facilitating detection.

### 3.3 Threshold of Detectability

Recall that the observed data is  $y_i = \mu x_i + \epsilon_i$   $i = 1, \dots, p$ , where  $\mu$  denotes the unknown signal strength,  $\mathbf{x}$  is the unknown activation pattern, and  $\epsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$ . The detection problem corresponds to the following hypothesis test:

$$H_0 : \mu = 0 \quad \text{vs.} \quad H_1 : \mu > 0$$

Projecting the network data onto the basis vectors  $\mathbf{b} \in \mathbb{B}$  yields the empirical transform coefficients  $\mathbf{b}_i^T \mathbf{y}$ . If the pattern  $\mathbf{x}$  is sparser in the transform domain, then its energy is concentrated in a few non-zero coefficients. Thus, the signal-to-noise ratio is boosted and detection is easier. To investigate the threshold of detectability for weak but structured activation patterns, we consider a simple test based on the maximum of the absolute values of the empirical transform coefficients  $\max_i |\mathbf{b}_i^T \mathbf{y}|$  as the test statistic. The following theorem provides an upper bound on the detection

threshold using the max statistic in the transform domain for patterns drawn from the tree-structured Ising model.

**Theorem 3.** *Consider a pattern  $\mathbf{x}$  drawn at random from a latent Ising model on a tree-structured graph with uniform degree  $d$  and depth  $L = \log_d p$ . Let  $\ell_0 = \frac{\alpha}{\beta}L$  and the interaction strength scales with the level  $\ell$  as  $\gamma_\ell = \ell\beta \log d$  for  $\ell \geq \ell_0$ , and  $\gamma_\ell = \infty$  for  $\ell < \ell_0$ .*

*With probability  $> 1 - 2\delta$  over the draw of the activation pattern, the test statistic  $\max_i |\mathbf{b}_i^T \mathbf{y}|$  drives the probability of false alarm and miss (conditioned on the draw of the pattern) to zero asymptotically as  $p \rightarrow \infty$  if the signal strength*

$$\mu > c \sqrt{2p^{-\kappa} \sigma^2 \log p},$$

where  $\kappa = \beta - \alpha > 0$  and  $c > 0$  is a constant.

Proof is given in the Appendix. We see that a polynomial improvement is attained if the activation pattern is sparser in a network transform domain. This is a significant improvement over canonical domain methods that do not exploit the structure of patterns and are limited to detecting signals with strength  $\mu > \sqrt{2\eta_\alpha \sigma^2 \log p}$  (where  $0 < \eta_\alpha < 1$  is independent of  $p$ ) (Ingster & Suslina 2002, Jin & Donoho 2004, Ingster et al. 2009).

## 4 Learning Clusters from Data

In practice, the pairwise similarities or covariances used for hierarchical clustering and constructing the proposed transform can only be estimated from data. Since the empirical covariance between network nodes can be learnt from multiple i.i.d. snapshots of network measurements, we now provide finite sample guarantees on the recovery of the multi-scale dependency structure from empirically estimated covariances. Analogous arguments can also be made for any similarity measure provided the empirical estimates satisfy a concentration inequality.

**Theorem 4.** *Consider noisy network measurements*

$$y_i = x_i + \epsilon_i \quad i = 1, \dots, p$$

*with known signal strength, where the noise variables  $\epsilon_i$  are independent  $\mathcal{N}(0, \sigma^2)$ . The  $x_i$  are independent of  $\epsilon_i$ , are uniformly bounded by  $M$ , and are zero-mean for simplicity. Assume  $\{\mathbb{E}[x_i x_j]\}$  satisfy the conditions of Lemma 1 for a hierarchical set of clusters  $\mathcal{H}^*$ . Let  $\tau$  denotes the smallest difference (gap) between the minimum pairwise covariance of leaf variables within any cluster and the maximum covariance between leaf variables in different clusters. Suppose we observe  $n$  i.i.d noisy realizations  $\{y_1^{(k)}, \dots, y_p^{(k)}\}_{k=1}^n$  of the  $p$  leaf variables, and  $\{\hat{r}_{ij} = \frac{1}{n} \sum_{k=1}^n y_i^{(k)} y_j^{(k)}\}$  denote the empirical covariances. Let  $\delta > 0$ . If  $n/\log n \geq \log(c_1 p^2 / \delta) / c_2 \tau^2$ , then with probability  $> 1 - \delta$ , the agglomerative clustering algorithm of Figure 1 applied to  $\{\hat{r}_{ij}\}$  recovers  $\mathcal{H}^*$ . Here  $c_1, c_2 > 0$  are constants that depend on  $M$  and  $\sigma^2$ .*




Figure 2: Performance comparison of global fusion, FDR, and the max statistic in transform and canonical domains, for weak patterns generated according to a hidden multi-scale Ising model.

The theorem implies that only  $O(\log p)$  measurements are needed to learn hierarchical clustering, and hence the proposed transform, in a network with  $p$  nodes.

## 5 Simulations

We simulated patterns from a multi-scale Ising model defined on a tree-structured graph with  $p = 1296$  leaf nodes with degree  $d = 6$  and depth  $L = 4$ . The network observations are modeled by adding white gaussian noise with standard deviation  $\sigma = 0.1$  to these patterns. This implies that a weak pattern is characterized by signal strength  $\mu < \sigma\sqrt{2\log p} = 0.38$ . We generate weak patterns with signal strength  $\mu$  varying from 0.06 to 0.2 and compare the detection performance of the max statistic in transform and canonical domains, and the global aggregate statistic, for a target false alarm probability of 0.05. We also compare to the FDR (False Discovery Rate) (Benjamini & Hochberg 1995) which is a canonical domain method that orders the measurements and thresholds them at a level that is adapted to the unknown sparsity level. The probability of detection as a function of signal strength is plotted in Figure 2. Detection in the transform domain clearly outperforms other methods since our construction exploits network node interactions.

The algorithmic complexity of hierarchical clustering  $p$  objects is  $O(p^2 \log p)$ , which essentially dominates the complexity of the detection procedure we propose.

## Appendix

**Proof of Theorem 1:** Each unbalanced Haar basis vector  $\mathbf{b} \in \mathbb{B}$  (except for the global summary vector  $\mathbf{1}_{\mathcal{L}}/\sqrt{|\mathcal{L}|}$ ) has one vanishing moment, i.e.  $\mathbf{b}^T \mathbf{1} = 0$ . Therefore, the only basis vectors with non-zero coefficients are the ones whose support contains a pair of nodes with different activation values. The number of node pairs with different activation values can be bounded by the total number of edge flips (variables that do not agree with their parent variables) in the tree. Let  $D_\ell$  denote the number of edge flips at level  $\ell$ . Since there are no more than  $dL$  basis vectors supported on a node pair with different activation values, the total number of non-zero coefficients  $\|\mathbb{B}^T \mathbf{x}\|_0 \leq dL \sum_\ell D_\ell$ .

Now observe that the tree-structured Ising model essentially specifies that edge flips are independent and occur with probability  $q_\ell = 1/(1 + e^{\gamma_\ell}) = 1/(1 + d^{\beta_\ell})$  at level  $\ell$ . That is, the number of flips per level  $D_\ell \sim \text{Binomial}(|E_\ell|, q_\ell)$  where  $E_\ell (= d^\ell)$  denotes the number of edges at level  $\ell$ . Let  $\ell' = L(1 - \beta) = (1 - \beta) \log_d p$ . Now  $d^{\ell(1-\beta)}/2 \leq |E_\ell|q_\ell \leq d^{\ell(1-\beta)}$ , and therefore  $|E_\ell|q_\ell \rightarrow \infty$  as  $p \rightarrow \infty$  for all  $\ell > \ell'$ . Invoking the relative Chernoff bound, we have: For any  $\ell > \ell'$ , with probability  $> 1 - \delta/L$ ,  $2^{-1}|E_\ell|q_\ell \leq D_\ell \leq 2|E_\ell|q_\ell$  for  $p$  large enough. We now have the following bound: With prob  $> 1 - \delta$

$$\begin{aligned} \|\mathbb{B}^T \mathbf{x}\|_0 &\leq dL \left( \sum_{\ell=1}^{\ell'} D_\ell + \sum_{\ell=\ell'+1}^L D_\ell \right) \\ &\leq dL \left( \sum_{\ell=1}^{\ell'} |E_\ell| + \sum_{\ell=\ell'+1}^L 2|E_\ell|q_\ell \right) \leq 3dL^2 d^{L(1-\beta)}. \end{aligned}$$

**Proof of Theorem 2:** For  $\ell < \ell_0$ ,  $\gamma_\ell = \infty$  implies that the probability of edge flip at level  $\ell$ ,  $q_\ell = 0$ . Following Theorem 1 proof, the transform sparsity bound still holds.

To evaluate the canonical domain sparsity, we condition on patterns for which the root variable is zero (inactive). Let  $A_\ell$  denote the number of variables that are active (take value 1) at level  $\ell$ . Since  $q_\ell = 0$  for  $\ell < \ell_0$ , there are no flips and hence no variables are active up to level  $\ell_0$ , i.e.  $A_\ell = 0$  for  $\ell < \ell_0$ . We argue that the canonical sparsity is governed by the number of nodes that are activated by flips at level  $\ell_0$ . Flips at lower levels might activate/de-activate some of the nodes but their effect is insignificant.

First, observe that the number of active variables at level  $\ell_0$ , conditioned on the root variable being inactive, is simply the number of edge flips  $D_{\ell_0}$  at level  $\ell_0$ , i.e.  $A_{\ell_0} = D_{\ell_0}$ . Consider  $\ell > \ell_0$ . Let  $M_\ell$  denote the number of active variables at level  $\ell$  whose parents were inactive, and let  $N_\ell$  denote the number of active variables at level  $\ell$  whose parents were also active. Therefore,  $A_\ell = M_\ell + N_\ell$ . Observe that, conditioned on the values of the variables at level  $\ell - 1$ ,

$$M_\ell | A_{\ell-1} \sim \text{Binomial}((|E_{\ell-1}| - A_{\ell-1})d, q_\ell)$$

$$N_\ell | A_{\ell-1} \sim \text{Binomial}(A_{\ell-1}d, 1 - q_\ell)$$

To gain some understanding for the canonical sparsity, we first look at the expected canonical sparsity. Note that  $\mathbb{E}[\|\mathbf{x}\|_0] = \mathbb{E}[A_L] = \mathbb{E}[\mathbb{E}[A_L|A_{L-1}]] = \mathbb{E}[\mathbb{E}[M_L + N_L|A_{L-1}]]$ . For the lower bound, we proceed as follows.  $\mathbb{E}[A_L] \geq \mathbb{E}[\mathbb{E}[N_L|A_{L-1}]] \geq \mathbb{E}[A_{L-1}]d(1 - q_L)$ . Now, repeatedly applying similar arguments for  $\ell > \ell_0$ ,

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}\|_0] &\geq \mathbb{E}[A_{\ell_0}]d^{L-\ell_0}\prod_{\ell>\ell_0}^L(1 - q_\ell) \\ &\geq |E_{\ell_0}|q_{\ell_0}d^{L-\ell_0}(1 - q_{\ell_0})^{L-\ell_0} \\ &\geq \frac{d^{\ell_0(1-\beta)}}{2}d^{L-\ell_0}(1 - d^{-\ell_0\beta})^{L-\ell_0} \\ &= \frac{1}{2}d^Ld^{-\ell_0\beta}(1 - p^{-\alpha})^{\log_d p^{1-\frac{\alpha}{\beta}}} \geq cd^Ld^{-\ell_0\beta} = cp^{1-\alpha}, \end{aligned}$$

where  $c < 1$ . The second step uses the fact that  $1 - q_\ell$  decreases with  $\ell$ , and that  $A_{\ell_0} = D_{\ell_0} \sim \text{Binomial}(|E_{\ell_0}|, q_{\ell_0})$ . The last inequality holds for large enough  $p$ . For the upper bound, we proceed as follows.

$$\begin{aligned} \mathbb{E}[A_L] &= \mathbb{E}[\mathbb{E}[M_L + N_L|A_{L-1}]] \\ &= \mathbb{E}[|E_{L-1}| - A_{L-1}dq_L + A_{L-1}d(1 - q_L)] \\ &\leq |E_{L-1}|dq_L + \mathbb{E}[A_{L-1}]d \end{aligned}$$

Repeatedly applying similar arguments for  $\ell > \ell_0$ , we get:

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}\|_0] &\leq \sum_{\ell=\ell_0}^{L-\ell_0} |E_{L-\ell}|d^\ell q_{L-\ell+1} + \mathbb{E}[A_{\ell_0}]d^{L-\ell_0} \\ &\leq \sum_{\ell=\ell_0}^{L-\ell_0} d^L d^{-(L-\ell+1)\beta} + |E_{\ell_0}|q_{\ell_0}d^{L-\ell_0} \\ &\leq Ld^L d^{-(\ell_0+1)\beta} + d^{\ell_0(1-\beta)}d^{L-\ell_0} \\ &\leq (L+1)d^L d^{-\ell_0\beta} \leq C(\log_d p)p^{1-\alpha}, \end{aligned}$$

where  $C > 1$ . The second step uses the fact that  $A_{\ell_0} = D_{\ell_0} \sim \text{Binomial}(|E_{\ell_0}|, q_{\ell_0})$ .

We now show that similar bounds on canonical sparsity hold with high probability as well. For this, we will invoke the relative Chernoff bound for binomial random variables  $M_\ell$  and  $N_\ell$ . First, we derive a lower bound on  $A_\ell$  for  $\ell > \ell_0$  recursively as follows. Recall that  $A_{\ell_0} = D_{\ell_0} \sim \text{Binomial}(|E_{\ell_0}|, q_{\ell_0})$  and using relative Chernoff bound as in the previous proof, w.p.  $> 1 - \delta/L$ ,  $A_{\ell_0} = D_{\ell_0} \geq \mathbb{E}[D_{\ell_0}]/2 = |E_{\ell_0}|q_{\ell_0}/2 \geq d^{\ell_0(1-\beta)}/4 \rightarrow \infty$  since  $\ell_0 = \frac{\alpha}{\beta}L = \frac{\alpha}{\beta}\log_d p \rightarrow \infty$ . Now  $A_{\ell_0+1} \geq N_{\ell_0+1}$ . And  $\mathbb{E}[N_{\ell_0+1}|A_{\ell_0}] = A_{\ell_0}d(1 - q_{\ell_0+1}) \geq A_{\ell_0}d(1 - q_{\ell_0}) \geq A_{\ell_0}d(1 - d^{-\ell_0\beta}) = A_{\ell_0}d(1 - p^{-\alpha})$ . Thus,  $\mathbb{E}[N_{\ell_0+1}|A_{\ell_0}] \rightarrow \infty$  w.p.  $> 1 - \delta/L$ . Conditioning on the values of the variables at level  $\ell_0$  and using relative Chernoff bound, we have with probability  $> 1 - 2\delta/L$ ,

$$\begin{aligned} A_{\ell_0+1} \geq N_{\ell_0+1} &\geq \mathbb{E}[N_{\ell_0+1}|A_{\ell_0}](1 - \epsilon_{\ell_0+1}) \\ &\geq A_{\ell_0}d(1 - p^{-\alpha})(1 - \epsilon_{\ell_0+1}) \\ \text{where } \epsilon_{\ell_0+1} &= \sqrt{\frac{3 \log(L/\delta)}{\mathbb{E}[N_{\ell_0+1}|A_{\ell_0}]}} \leq \sqrt{\frac{3 \log(L/\delta)}{A_{\ell_0}d(1 - p^{-\alpha})}} \\ &\leq c' p^{-\frac{\alpha}{2\beta}(1-\beta)} \sqrt{\log \log p} < 1 \end{aligned}$$

for  $p$  large enough and  $c' > 0$  is a constant. Notice that  $A_{\ell_0+1} \rightarrow \infty$  with probability  $> 1 - 2\delta/L$ . Now consider any  $\ell > \ell_0$  and assume that for all  $\ell \geq \ell' > \ell_0$ ,  $A_{\ell'} \geq A_{\ell'-1}d(1 - p^{-\alpha})(1 - \epsilon_{\ell'})$ , where  $\epsilon_{\ell'} \leq c' p^{-\frac{\alpha}{2\beta}(1-\beta)} \sqrt{\log \log p} < 1$ , and  $A_{\ell'} \rightarrow \infty$  with probability  $> 1 - (\ell' - \ell_0 + 1)\delta/L$ . We show that similar arguments are true for  $A_{\ell+1}$ . Recall that  $A_{\ell+1} \geq N_{\ell+1}$ . And  $\mathbb{E}[N_{\ell+1}|A_\ell] = A_\ell d(1 - q_{\ell+1}) \geq A_\ell d(1 - q_{\ell_0}) \geq A_\ell d(1 - p^{-\alpha})$ . Thus,  $\mathbb{E}[N_{\ell+1}|A_\ell] \rightarrow \infty$  w.h.p. since  $A_\ell \rightarrow \infty$ . Now, conditioning on the values of the variables at level  $\ell$  and using relative Chernoff bound, we have with probability  $> 1 - (\ell - \ell_0 + 2)\delta/L$ ,

$$\begin{aligned} A_{\ell+1} \geq N_{\ell+1} &\geq \mathbb{E}[N_{\ell+1}|A_\ell](1 - \epsilon_{\ell+1}) \\ &\geq A_\ell d(1 - p^{-\alpha})(1 - \epsilon_{\ell+1}) \end{aligned}$$

$$\begin{aligned} \text{where } \epsilon_{\ell+1} &= \sqrt{\frac{3 \log(L/\delta)}{\mathbb{E}[N_{\ell+1}|A_\ell]}} \leq \sqrt{\frac{3 \log(L/\delta)}{A_\ell d(1 - p^{-\alpha})}} \\ &\leq \sqrt{\frac{3 \log(L/\delta)}{A_{\ell_0}(1 - p^{-\alpha})^{\ell+1-\ell_0} d^{\ell+1-\ell_0} \prod_{\ell'=\ell_0+1}^\ell (1 - \epsilon_{\ell'})}} \\ &\leq c' p^{-\frac{\alpha}{2\beta}(1-\beta)} \sqrt{\log \log p} \end{aligned}$$

The last step follows by recalling that  $A_{\ell_0} \geq d^{\ell_0(1-\beta)}/4 = p^{\frac{\alpha}{\beta}(1-\beta)}/4$  and  $(1 - p^{-\alpha})^{\ell+1-\ell_0} \geq (1 - p^{-\alpha})^{L+1-\ell_0} = (1 - p^{-\alpha})^{(1-\alpha/\beta)\log_d p+1} > c'$  for large enough  $p$ . Also,  $\epsilon_{\ell'} \leq 1/2$  for large enough  $p$  and hence  $d^{\ell+1-\ell_0} \prod_{\ell'=\ell_0+1}^\ell (1 - \epsilon_{\ell'}) \geq d(d/2)^{\ell-\ell_0} \geq 1$ . Thus we get, with probability  $> 1 - \delta$ , for all  $\ell > \ell_0$

$$A_\ell \geq A_{\ell_0}d^{\ell-\ell_0}(1 - p^{-\alpha})^{\ell-\ell_0} \prod_{\ell'=\ell_0+1}^\ell (1 - \epsilon_{\ell'})$$

where  $\epsilon_{\ell'} \leq c' p^{-\frac{\alpha}{2\beta}(1-\beta)} \sqrt{\log \log p} < 1$ . Finally, we have a lower bound on the canonical sparsity as follows: With probability  $> 1 - \delta$ ,  $\|\mathbf{x}\|_0 = A_L$  and

$$\begin{aligned} A_L &\geq A_{\ell_0}d^{L-\ell_0}((1 - p^{-\alpha})(1 - c' p^{-\frac{\alpha(1-\beta)}{2\beta}} \log p))^{L-\ell_0} \\ &\geq cd^{\ell_0(1-\beta)}d^{L-\ell_0} = cd^L d^{-\ell_0\beta} = cp^{1-\alpha} \end{aligned}$$

where we use the fact that  $(1 - p^{-\alpha})^{\log_d p^b} \geq c > 0$  for large enough  $p$ . Also note that  $c < 1$ .

For the upper bound, recall that  $A_\ell = M_\ell + N_\ell$  and proceed recursively using relative Chernoff bound for both  $M_\ell, N_\ell$ . For details, see (Singh, Nowak & Calderbank 2010).

**Proof of Theorem 3:** Consider the threshold  $t = \sqrt{2\sigma^2(1+c)\log p}$ , where  $c > 0$  is an arbitrary constant. Since the proposed transform is orthonormal, it is easy to see that under the null hypothesis  $H_0$  (no activation), the empirical transform coefficients  $\mathbf{b}_i^T \mathbf{y} \sim \mathcal{N}(0, \sigma^2)$ . Therefore, the false alarm probability can be bounded as follows:

$$\begin{aligned} P_{H_0}(\max_i |\mathbf{b}_i^T \mathbf{y}| > t) &= 1 - \prod_{i=1}^p P_{H_0}(|\mathbf{b}_i^T \mathbf{y}| \leq t) \\ &\leq 1 - (1 - 2e^{-t^2/2\sigma^2})^p = 1 - \left(1 - \frac{1}{p^{1+c}}\right)^p \rightarrow 0 \end{aligned}$$

Under the alternate hypothesis  $H_1$ , the empirical transform coefficients  $\mathbf{b}_i^T \mathbf{y} \sim \mathcal{N}(\mu \mathbf{b}_i^T \mathbf{x}, \sigma^2)$ . Therefore, the miss probability can be bounded as follows:

$$\begin{aligned} P_{H_1}(\max_i |\mathbf{b}_i^T \mathbf{y}| \leq t) &\leq \Pi_{i: \mathbf{b}_i^T \mathbf{x} \neq 0} P(|\mathcal{N}(\mu \mathbf{b}_i^T \mathbf{x}, \sigma^2)| \leq t) \\ &= \Pi_{i: \mathbf{b}_i^T \mathbf{x} > 0} P(\mathcal{N}(0, \sigma^2) \leq t - \mu |\mathbf{b}_i^T \mathbf{x}|) \\ &\quad \cdot \Pi_{i: \mathbf{b}_i^T \mathbf{x} < 0} P(\mathcal{N}(0, \sigma^2) \geq -t + \mu |\mathbf{b}_i^T \mathbf{x}|) \\ &= \Pi_{i: \mathbf{b}_i^T \mathbf{x} \neq 0} P(\mathcal{N}(0, \sigma^2) \leq t - \mu |\mathbf{b}_i^T \mathbf{x}|) \\ &\leq P(\mathcal{N}(0, \sigma^2) \leq t - \mu \max_i |\mathbf{b}_i^T \mathbf{x}|) \end{aligned}$$

In the second step we use the fact that  $P(|a| \leq t) \leq P(a \leq t)$  and  $P(|a| \leq t) \leq P(a \geq -t)$ . Thus, the miss probability goes to zero if  $\mu \max_i |\mathbf{b}_i^T \mathbf{x}| > (1 + c')t$  for any  $c' > 0$ .

The detectability threshold follows by deriving a lower bound for the largest absolute transform coefficient. The energy in the largest transform coefficient is at least as large as the average energy per non-zero coefficient, and hence  $\max_i |\mathbf{b}_i^T \mathbf{x}| \geq \sqrt{\|\mathbf{x}\|_0 / \|\mathbb{B}^T \mathbf{x}\|_0}$ . Invoking Theorem 2 for patterns that correspond to the root value zero, with probability  $> 1 - 2\delta$ ,  $\max_i |\mathbf{b}_i^T \mathbf{x}| \geq c p^{(\beta-\alpha)/2}$ , where  $c > 0$  is a constant. Patterns that correspond to the root variable taking value one are canonically non-sparse with even larger  $\|\mathbf{x}\|_0$ . Thus, the same lower bound holds in this case.

**Proof of Theorem 4:** The true hierarchical structure  $\mathcal{H}^*$  between the leaf variables can be recovered if the empirical covariances  $\{\hat{r}_{ij}\}$  satisfy the conditions of Lemma 1. Since the true covariance  $r_{ij} = \mathbb{E}[(y_i y_j)] = \mathbb{E}[(x_i x_j)]$  for  $i \neq j$  (the auto-covariances are not important for clustering) satisfy these conditions, a sufficient condition for recovery of the structure is that the deviation between true and empirical covariance of the observed variables is less than  $\tau/2$ , i.e.  $\max_{(i,j)} |\hat{r}_{ij} - r_{ij}| < \tau/2$ .

To establish this, we study the concentration of the empirical covariances around the true covariances. Since  $\mathbf{y}$  is the sum of a bounded random variable and gaussian noise, it can be shown (Singh et al. 2010) that  $v_k := y_i^{(k)} y_j^{(k)}$  satisfies the moment condition  $\mathbb{E}[|v_k - \mathbb{E}[v_k]|^p] \leq p! \text{var}(v_k) h^{p-2}/2$  for integers  $p \geq 2$  and some constant  $h > 0$ . We can now invoke the Bernstein inequality:

$$P\left(\frac{1}{n} \sum_{k=1}^n (v_k - \mathbb{E}[v_k]) > \frac{2t}{n} \sqrt{\sum_{k=1}^n \text{var}(v_k)}\right) < e^{-t^2}$$

for  $0 < t \leq \sqrt{\sum_{k=1}^n \text{var}(v_k)}/(2h)$ . Now, straight-forward computations show that  $c_1 := \sigma^4 \leq \text{var}(v_k) \leq 2\sigma^4 + 4M^2\sigma^2 + 4M^4 =: c_2$ . And we get

$$P\left(\frac{1}{n} \sum_{k=1}^n (v_k - \mathbb{E}[v_k]) > \frac{2t\sqrt{c_2}}{\sqrt{n}}\right) < e^{-t^2}$$

Let  $t = \sqrt{n}\tau/(4\sqrt{c_2} \log n)$ . Then  $0 < t \leq \sqrt{nc_1}/(2h) \leq \sqrt{\sum_{k=1}^n \text{var}(v_k)}/(2h)$  for large enough  $n$  and hence  $t$  satisfies the desired conditions. Similar arguments show that

$-v_k$  also satisfies the moment condition. Using this and taking union bound over all elements in similarity matrix,

$$P(\max_{ij} |\hat{r}_{ij} - r_{ij}| > \tau/2) < 2p^2 e^{-n\tau^2/(16c_2 \log n)}.$$

Thus, the covariance clustering algorithm of Figure 1 recovers  $\mathcal{H}^*$  with probability  $> 1 - \delta$  if  $n/\log n \geq 16c_2 \log(2p^2/\delta)/\tau^2$ .

**Acknowledgements:** This work was supported in part by NSF under grant DMS 0701226, by ONR under grant N00173-06-1-G006, and by AFOSR under grants FA9550-05-1-0443 and FA9550-09-1-0423.

## References

- A.-Berry, L., Broutin, N., Devroye, L. & Lugosi, G. (2009), ‘On combinatorial testing problems, <http://arxiv.org/abs/0908.3437>’.
- A.-Castro, E., Candés, E. J. & Durand, A. (2010), ‘Detection of an abnormal cluster in a network, <http://arxiv.org/abs/1001.3209>’.
- A.-Castro, E., Candés, E. J., Helgason, H. & Zeitouni, O. (2007), ‘Searching for a trail of evidence in a maze’, *Annals of Statistics* **36**, 1726–1757.
- A.-Castro, E., Donoho, D. L. & Huo, X. (2005), ‘Near-optimal detection of geometric objects by fast multiscale methods’, *IEEE Trans. Info. Theory* **51**(7), 2402–2425.
- Benjamini, Y. & Hochberg, Y. (1995), ‘Controlling the false discovery rate: a practical and powerful approach to multiple testing’, *Journal Royal Stat. Soc. Series B* **57**, 289–300.
- Falk, H. (1975), ‘Ising spin system on a cayley tree: Correlation decomposition and phase transition’, *Physical Review B* **12**(11).
- Girvan, M. & Newman, M. E. J. (2002), ‘Community structure in social and biological networks’, *Proc. Natl. Acad. Sci. USA* **99**(12), 7821–7826.
- Ingster, Y. I., Pouet, C. & Tsybakov, A. B. (2009), ‘Sparse classification boundaries, <http://arxiv.org/abs/0903.4807>’.
- Ingster, Y. I. & Suslina, I. A. (2002), *Nonparametric goodness-of-fit testing under Gaussian models*.
- Jager, L. & Wellner, J. A. (2007), ‘Goodness-of-fit tests via phi-divergences’, *Annals of Statistics* **35**, 2018–2053.
- Jin, J. & Donoho, D. L. (2004), ‘Higher criticism for detecting sparse heterogeneous mixtures’, *Annals of Statistics* **32**(3), 962–994.
- Lee, A. B., Nadler, B. & Wasserman, L. (2008), ‘Treelets - an adaptive multi-scale basis for sparse unordered data’, *Annals of Applied Statistics* **2**(2), 435–471.
- Murtagh, F. (2007), ‘The haar wavelet transform of a dendrogram’, *J. Classification* **24**, 3–32.
- Ramasubramanian, R., Malkhi, D., Kuhn, F., Balakrishnan, M. & Akella, A. (2009), On the treeness of internet latency and bandwidth, in ‘Proc. of SIGMETRICS, Seattle, WA’.
- Sankaranarayanan, L., Kramer, G. & Mandayam, N. B. (2004), Hierarchical sensor networks: capacity bounds and cooperative strategies using the multiple-access relay channel model, in ‘IEEE Comm. Soc. Conf. Sensor & Ad Hoc Comm. Nwks’.
- Singh, A., Nowak, R. & Calderbank, R. (2010), ‘Detecting weak but hierarchically-structured patterns in networks, <http://arxiv.org/abs/1003.0205v1>’.
- Varshney, P. K. (1996), *Distributed Detection and Data Fusion*, Springer-Verlag New York Inc.
- Yu, H. & Gerstein, M. (2006), ‘Genomic analysis of the hierarchical structure of regulatory networks’, *Proc. Natl. Acad. Sci. USA* **103**, 14724–14731.