



HAL
open science

Fast Randomized Semi-Supervised Clustering

Alaa Saade, Florent Krzakala, Marc Lelarge, Lenka Zdeborová

► **To cite this version:**

Alaa Saade, Florent Krzakala, Marc Lelarge, Lenka Zdeborová. Fast Randomized Semi-Supervised Clustering. *Journal of Physics: Conference Series*, 2018, 1036, pp.012015. 10.1088/1742-6596/1036/1/012015 . cea-04542690

HAL Id: cea-04542690

<https://cea.hal.science/cea-04542690v1>

Submitted on 11 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

PAPER • OPEN ACCESS

Fast Randomized Semi-Supervised Clustering

To cite this article: Alaa Saade *et al* 2018 *J. Phys.: Conf. Ser.* **1036** 012015

View the [article online](#) for updates and enhancements.

You may also like

- [A novel task-oriented optimal design for P300-based brain-computer interfaces](#)
Zongtan Zhou, Erwei Yin, Yang Liu et al.
- [Weighted community detection and data clustering using message passing](#)
Cheng Shi, Yanchen Liu and Pan Zhang
- [Diffusion-based location-aware recommender systems](#)
Hao Liao, Xiaojie Zhang, Zhongtian Long et al.

PRIME
PACIFIC RIM MEETING
ON ELECTROCHEMICAL
AND SOLID STATE SCIENCE

HONOLULU, HI
Oct 6–11, 2024

Abstract submission deadline:
April 12, 2024

Learn more and submit!

Joint Meeting of

The Electrochemical Society
•
The Electrochemical Society of Japan
•
Korea Electrochemical Society

Fast Randomized Semi-Supervised Clustering

Alaa Saade¹, Florent Krzakala^{1,2}, Marc Lelarge³, Lenka Zdeborová⁴

¹ Laboratoire de Physique Statistique (CNRS UMR-8550), PSL Universités and École Normale Supérieure, 75005 Paris (e-mail: alaa.saade,florent.krzakala@ens.fr)

² Sorbonne Universités, UPMC Univ. Paris 06

³ INRIA and École Normale Supérieure Paris, France (e-mail: marc.lelarge@ens.fr)

⁴ Institut de Physique Théorique CEA Saclay and CNRS, France (e-mail: lenka.zdeborova@gmail.com)

Abstract. We consider the problem of clustering partially labeled data from a minimal number of randomly chosen pairwise comparisons between the items. We introduce an efficient local algorithm based on a power iteration of the non-backtracking operator and study its performance on a generative model. For the case of two clusters, we give bounds on the classification error and show that a small error can be achieved from $O(n)$ randomly chosen measurements, where n is the number of items in the dataset. Our algorithm is therefore efficient both in terms of time and space complexities. We also investigate numerically the performance of the algorithm on synthetic and real-world data.

1. Introduction

Similarity-based clustering aims at classifying data points into homogeneous groups based on some measure of their resemblance. The problem can be stated formally as follows: given n items $\{x_i\}_{i \in [n]} \in \mathcal{X}^n$, and a symmetric similarity function $s : \mathcal{X}^2 \rightarrow \mathbb{R}$, the aim is to cluster the dataset from the knowledge of the pairwise similarities $s_{ij} := s(x_i, x_j)$, for $1 \leq i < j \leq n$. This information is usually represented in the form of a weighted similarity graph $G = (V, E)$ where the nodes represent the items of the dataset and the edges are weighted by the pairwise similarities. Popular choices for the similarity graph are the fully-connected graph, or the k -nearest neighbors graph, for a suitable k (see e.g. [1] for a review). Both choices, however, require the computation of a large number of pairwise similarities, typically $O(n^2)$. For large datasets, with n in the millions or billions, or large-dimensional data, where computing each similarity s_{ij} is costly, the complexity of this procedure is often prohibitive, both in terms of computational and memory requirements.

It is then natural to ask whether it is really required to compute as many as $O(n^2)$ similarities to accurately cluster the data. In the absence of additional information on the data, a reasonable alternative is to compare each item to a small number of other items in the dataset, chosen uniformly at random. Random subsampling methods are a well-known means of reducing the complexity of a problem, and they have been shown to yield substantial speed-ups in clustering [2] and low-rank approximation [3, 4]. In particular, [5] recently showed that an unsupervised spectral method based on the principal eigenvectors of the non-backtracking operator of [6] can cluster the data better than chance from only $O(n)$ similarities.

In this paper, we build upon previous work by considering two variations motivated by real-world applications. The first question we address is how to incorporate the knowledge of the



labels of a small fraction of the items to aid clustering of the whole dataset, resulting in a more efficient algorithm. This question, referred to as semi-supervised clustering, is of broad practical interest [7, 8]. For instance, in a social network, we may have pre-identified individuals of interest, and we might be looking for other individuals sharing similar characteristics. In biological networks, the function of some genes or proteins may have been determined by costly experiments, and we might seek other genes or proteins sharing the same function. More generally, efficient human-powered methods such as crowdsourcing can be used to accurately label part of the data [9, 10], and we might want to use this knowledge to cluster the rest of the dataset at no additional cost.

The second question we address is the number of randomly chosen pairwise similarities that are needed to achieve a given classification error. Previous work has mainly focused on two related, but different questions. One line of research has been interested in exact recovery, i.e. how many measurements are necessary to exactly cluster the data. Note that for exact recovery to be possible, it is necessary to choose at least $O(n \log n)$ random measurements for the similarity graph to be connected with high probability. On simple models, [11, 12, 13] showed that this scaling is also sufficient for exact recovery. On the sparser end of the spectrum, [14, 15, 16, 5] have focused on the detectability threshold, i.e. how many measurements are needed to cluster the data better than chance. On simple models, this threshold is typically achievable with $O(n)$ measurements only. While this scaling is certainly attractive for large problems, it is important for practical applications to understand how the expected classification error decays with the number of measurements.

To answer these two questions, we introduce a highly efficient, local algorithm based on a power iteration of the non-backtracking operator. For the case of two clusters, we show on a generative model that the classification error decays exponentially with the number of measured pairwise similarities, thus allowing the algorithm to cluster data to arbitrary accuracy while being efficient both in terms of time and space complexities. We demonstrate the good performance of this algorithm on both synthetic and real-world data, and compare it to the popular label propagation algorithm [8].

2. Algorithm and guarantee

2.1. Algorithm for 2 clusters

Consider n items $\{x_i\}_{i \in [n]} \in \mathcal{X}^n$ and a symmetric similarity function $s : \mathcal{X}^2 \rightarrow \mathbb{R}$. The choice of the similarity function is problem-dependent, and we will assume one has been chosen. For concreteness, s can be thought of as a decreasing function of a distance if \mathcal{X} is an Euclidean space. The following analysis, however, applies to a generic function s , and our bounds will depend explicitly on its statistical properties. We assume that the true labels $(\sigma_i = \pm 1)_{i \in \mathcal{L}}$ of a subset $\mathcal{L} \subset [n]$ of items is known. Our aim is to find an estimate $(\hat{\sigma}_i)_{i \in [n]}$ of the labels of all the items, using a small number of similarities. More precisely, let E be a random subset of all the $\binom{n}{2}$ possible pairs of items, containing each given pair $(ij) \in [n]^2$ with probability α/n , for some $\alpha > 0$. We compute only the $O(\alpha n)$ similarities $(s_{ij} := s(x_i, x_j))_{(ij) \in E}$ of the pairs thus chosen.

From these similarities, we define a weighted similarity graph $G = (V, E)$ where the vertices $V = [n]$ represent the items, and each edge $(ij) \in E$ carries a weight $w_{ij} := w(s_{ij})$, where w is a weighting function. Once more, we will consider a generic function w in our analysis, and discuss the performance of our algorithm as a function of the choice of w . In particular, we show in section 2.2 that there is an optimal choice of w when the data is generated from the symmetric model defined in (2) below. However, in practice, the main purpose of w is to center the similarities, i.e. we will take in our numerical simulations $w(s) = s - \bar{s}$, where \bar{s} is the empirical mean of the observed similarities. The necessity to center the similarities is discussed in the following. Note that the graph G is a weighted version of an Erdős-Rényi random graph with average degree α , which controls the sampling rate: a larger α means more

pairwise similarities are computed, at the expense of an increase in complexity. Algorithm 1 describes our clustering procedure for the case of 2 clusters. We denote by ∂i the set of neighbors of node i in the graph G , and by \vec{E} the set of directed edges of G .

Algorithm 1 Non-backtracking local walk (2 clusters)

Input: $n, \mathcal{L}, (\sigma_i = \pm 1)_{i \in \mathcal{L}}, E, (w_{ij})_{(ij) \in E}, k_{\max}$

Output: Cluster assignments $(\hat{\sigma}_i)_{i \in [n]}$

- 1: **Initialize** the messages $v^{(0)} = (v_{i \rightarrow j}^{(0)})_{(i \rightarrow j) \in \vec{E}}$
 - 2: **for all** $(i \rightarrow j) \in \vec{E}$ **do**
 - 3: **if** $i \in \mathcal{L}$ **then** $v_{i \rightarrow j}^{(0)} \leftarrow \sigma_i$
 - 4: **else** $v_{i \rightarrow j}^{(0)} \leftarrow \pm 1$ uniformly at random
 - 5: **Iterate** for $k = 1, \dots, k_{\max}$
 - 6: **for** $(i \rightarrow j) \in \vec{E}$ **do** $v_{i \rightarrow j}^{(k)} \leftarrow \sum_{l \in \partial i \setminus j} w_{il} v_{l \rightarrow i}^{(k-1)}$
 - 7: **Pool** the messages
 - 8: **for** $i \in [n]$ **do** $\hat{v}_i \leftarrow \sum_{l \in \partial i} w_{il} v_{l \rightarrow i}^{(k_{\max})}$
 - 9: **Output** the assignments
 - 10: **for** $i \in [n]$ **do** $\hat{\sigma}_i \leftarrow \text{sign}(\hat{v}_i)$
-

This algorithm can be thought of as a linearized version of a belief propagation algorithm, that iteratively updates messages on the directed edges of the similarity graph, by assigning to each message the weighted sum of its incoming messages. More precisely, Algorithm 1 can be observed to approximate the leading eigenvector of the non-backtracking operator B , whose elements are defined, for $(i \rightarrow j), (k \rightarrow l) \in \vec{E}$, by

$$B_{(i \rightarrow j), (k \rightarrow l)} := w_{kl} \mathbf{1}(i = l) \mathbf{1}(k \neq j). \quad (1)$$

It is therefore close in spirit to the unsupervised spectral methods introduced by [6, 16, 5], which rely on the computation of the principal eigenvector of B . On sparse graphs, methods based on the non-backtracking operator are known to perform better than traditional spectral algorithms based e.g. on the adjacency matrix, or random walk matrix of the sparse similarity graph, which suffer from the appearance of large eigenvalues with localized eigenvectors (see e.g. [6, 17]). In particular, we will see on the numerical experiments of Section 3 that Algorithm 1 outperforms the label propagation algorithm, based on an iteration of the random walk matrix.

However, in contrast with the past spectral approaches based on B , our algorithm is local, in that the estimate $\hat{\sigma}_i$ for a given item $i \in [n]$ depends only on the messages on the edges that are at most $k_{\max} + 1$ steps away from i in the graph G . This fact will prove essential in our analysis. Indeed, we will show that in our semi-supervised setting, a finite number of iterations (independent of n) is enough to ensure a low classification error. On the other hand, in the unsupervised setting, we expect local algorithms not to be able to find large clusters in a graph, a limitation that has already been highlighted on the related problems of finding large independent sets on graphs [18] and community detection [19]. On the practical side, the local nature of Algorithm 1 leads to a drastic improvement in running time. Indeed, in order to compute the leading eigenvector of B , a number of iterations k scaling with the number n of items is required [20]. Here, on the contrary, the number of iterations stays independent of the size n of the dataset.

2.2. Model and guarantee

To evaluate the performance of Algorithm 1, we consider the following semi-supervised variant of the labeled stochastic block model [15]. This model extends the popular stochastic block model by allowing generic edge-weights and is used as a benchmark for a theoretical analysis of graph clustering [12]. We assign n items to 2 predefined clusters of the same average size $n/2$, by drawing for each item $i \in [n]$ a cluster label $\sigma_i \in \{\pm 1\}$ with uniform probability $1/2$. We choose uniformly at random ηn items to form a subset $\mathcal{L} \subset [n]$ of items whose label is revealed, so that η is the fraction of labeled data. Next, we choose which pairs of items will be compared by generating an Erdős-Rényi random graph $G = (V = [n], E) \in \mathcal{G}(n, \alpha/n)$, for some constant $\alpha > 0$, independent of n . We will assume that the similarity s_{ij} between two items i and j is a random variable depending only on the labels of the items i and j . More precisely, we consider the symmetric model

$$\mathbb{P}(s_{ij} = s | \sigma_i, \sigma_j) = \begin{cases} p_{\text{in}}(s) & \text{if } \sigma_i = \sigma_j, \\ p_{\text{out}}(s) & \text{if } \sigma_i \neq \sigma_j, \end{cases} \quad (2)$$

where p_{in} (resp. p_{out}) is the distribution of the similarities between items within the same cluster (resp. different clusters). The properties of the weighting function w will determine the performance of our algorithm through the two following quantities. Define $2\Delta(w) := \mathbb{E}[w_{ij} | \sigma_j = \sigma_j] - \mathbb{E}[w_{ij} | \sigma_j \neq \sigma_j]$, the difference in expectation between the weights inside a cluster and between different clusters. Define also $\Sigma(w)^2 := \mathbb{E}[w^2]$, the second moment of the weights. Our first result (proved in section 4) is concerned with what value of α is required to improve the initial labeling with Algorithm 1.

Theorem 1. *Assume a similarity graph G with n items and a labeled set \mathcal{L} of size ηn to be generated from the symmetric model (2) with 2 clusters. Define $\tau(\alpha, w) := \frac{\alpha \Delta(w)^2}{\Sigma(w)^2}$. If $\Delta(w) > 0$, then there exists a constant $C > 0$ such that the estimates $\hat{\sigma}_i$ from k iterations of Algorithm 1 achieve*

$$\frac{1}{n} \sum_{i=1}^n \mathbb{P}(\sigma_i \neq \hat{\sigma}_i) \leq 1 - r_{k+1} + C \frac{\alpha^{k+1} \log n}{\sqrt{n}}, \quad (3)$$

where $r_0 = \eta^2$ and for $0 \leq l \leq k$,

$$r_{l+1} = \frac{\tau(\alpha, w) r_l}{1 + \tau(\alpha, w) r_l}. \quad (4)$$

To understand the content of this bound, we consider the limit of a large number of items $n \rightarrow \infty$, so that the last term of (3) vanishes. Note first that if $\tau(\alpha, w) > 1$, then starting from any positive initial condition, r_k converges to $(\tau(\alpha, w) - 1)/\tau(\alpha, w) > 0$ in the limit where the number of iterations $k \rightarrow \infty$. A random guess on the unlabeled points yields an asymptotic error of $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{P}(\sigma_i \neq \hat{\sigma}_i) = (1 - \eta)/2$, so that a sufficient condition for Algorithm 1 to improve the initial partial labeling, after a certain number of iterations $k(\tau(\alpha, w), \eta)$ independent of n , is

$$\tau(\alpha, w) > \frac{2}{1 - \eta}. \quad (5)$$

Note that there is an optimal choice of weighting function w which maximizes $\tau(\alpha, w)$, namely (consistently with [14]):

$$\begin{aligned} w^*(s) &:= \frac{p_{\text{in}}(s) - p_{\text{out}}(s)}{p_{\text{in}}(s) + p_{\text{out}}(s)} \\ \iff \tau(\alpha, w^*) &= \frac{\alpha}{2} \int ds \frac{(p_{\text{in}}(s) - p_{\text{out}}(s))^2}{p_{\text{in}}(s) + p_{\text{out}}(s)}, \end{aligned} \quad (6)$$

which, however, requires knowing the parameters of the model. It is then informative to compare our bound (5) to known optimal asymptotic bounds in the unsupervised setting. In the limit of vanishing supervision $\eta \rightarrow 0$, the bound (5) guarantees improving the initial labeling if $\tau(\alpha, w^*) > 2 + O(\eta)$. In the unsupervised setting ($\eta = 0$), it has been shown by [14] that if $\tau(\alpha, w^*) < 1$, no algorithm, either local or global, can cluster the graph better than random guessing. If $\tau(\alpha, w^*) > 1$, on the other hand, then a global spectral method based on the principal eigenvectors of the non-backtracking operator improves over random guessing [5]. This suggests that, in the limit of vanishing supervision $\eta \rightarrow 0$, the bound (5) is close to optimal, but off by a factor of 2.

Note, however, that Theorem 1 applies to a generic weighting function w . In particular, while the optimal choice (6) is not practical, theorem 1 guarantees that Algorithm 1 retains the ability to improve the initial labeling from a small number of measurements, as long as $\Delta(w) > 0$. With the choice $w(s) = s - \bar{s}$ advocated for in section 2.1, we have $2\Delta(w) = \mathbb{E}[s_{ij}|\sigma_j = \sigma_j] - \mathbb{E}[s_{ij}|\sigma_j \neq \sigma_j]$. Therefore Algorithm 1 improves over random guessing for α large enough if the similarity between items in the same cluster is larger in expectation than the similarity between items in different clusters, which is a reasonable requirement. Note that the hypotheses of Theorem 1 do not require the weighting function w to be centered. However, it is easy to check that if $\mathbb{E}[w] \neq 0$, defining a new weighting function by $w'(s) := w(s) - \mathbb{E}[w]$, we have $\tau(\alpha, w') > \tau(\alpha, w)$, so that the bound (3) is improved.

While Theorem 1 guarantees improving the initial clustering from a small sampling rate α , it provides a rather loose bound on the expected error when α becomes larger. The next theorem addresses this regime. A proof is given in Section 5.

Theorem 2. *Assume a similarity graph G with n items and a labeled set \mathcal{L} of size ηn to be generated from the symmetric model (2) with 2 clusters. Assume further that the weighting function w is bounded: $\forall s \in \mathbb{R}, |w(s)| \leq 1$. Define $\tau(\alpha, w) := \frac{\alpha \Delta(w)^2}{\Sigma(w)^2}$. If $\alpha \Delta(w) > 1$ and $\alpha \Sigma(w)^2 > 1$, then there exists a constant $C > 0$ such that the estimates $\hat{\sigma}_i$ from k iterations of Algorithm 1 achieve*

$$\frac{1}{n} \sum_{i=1}^n \mathbb{P}(\sigma_i \neq \hat{\sigma}_i) \leq \exp \left[-\frac{q_{k+1}}{4} \min \left(1, \frac{\Sigma(w)^2}{\Delta(w)} \right) \right] + C \frac{\alpha^{k+1} \log n}{\sqrt{n}}, \quad (7)$$

where $q_0 = 2\eta^2$ and for $0 \leq l \leq k$,

$$q_{l+1} = \frac{\tau(\alpha, w) q_l}{1 + 3/2 \max(1, q_l)}. \quad (8)$$

Note that by linearity of Algorithm 1, the condition $\forall s, |w(s)| \leq 1$ can be relaxed to w bounded. It is once more instructive to consider the limit of large number of items $n \rightarrow \infty$. Starting from any initial condition, if $\tau(\alpha, w) < 5/2$, then $q_k \xrightarrow[k \rightarrow \infty]{} 0$ so that the bound (7) is uninformative. On the other hand, if $\tau(\alpha, w) > 5/2$, then starting from any positive initial condition, $q_k \xrightarrow[k \rightarrow \infty]{} \frac{2}{3}(\tau(\alpha, w) - 1) > 0$. This bound therefore shows that on a model with a given distribution of similarities (2) and a given weighting function w , an error smaller than ϵ can be achieved from $\alpha n = O(n \log 1/\epsilon)$ measurements, in the limit $\epsilon \rightarrow 0$, with a finite number of iterations $k(\tau(\alpha, w), \eta, \epsilon)$ independent of n . We note that this result is the analog, for a weighted graph, of the recent results of [21] who show that in the stochastic block model, a local algorithm similar to Algorithm 1 achieves an error decaying exponentially as a function of a relevant signal to noise ratio.

2.3. More than 2 clusters

Algorithm 2 gives a natural extension of our algorithm to $q > 2$ clusters. In this case, we expect the non-backtracking operator B defined in equation (1) to have $q - 1$ large eigenvalues, with eigenvectors correlated with the types σ_i of the items (see [5]). We use a deflation-based power iteration method [22] to approximate these eigenvectors, starting from informative initial conditions incorporating the knowledge drawn from the partially labeled data. Numerical simulations illustrating the performance of this algorithm are presented in Section 3. Note that each deflated matrix B_c for $c \geq 2$ is a rank- $(c - 1)$ perturbation of a sparse matrix, so that the power iteration can be done efficiently using sparse linear algebra routines. In particular, both Algorithms 1 and 2 have a time and space complexities linear in the number of items n .

Algorithm 2 Non-backtracking local walk (q clusters)

Input: $n, q, \mathcal{L}, (\sigma_i \in [q])_{i \in \mathcal{L}}, E, (w_{ij})_{(ij) \in E}, k_{\max}$

Output: Cluster assignments $(\hat{\sigma}_i)_{i \in [n]}$

- 1: $B_1 \leftarrow B$ where B is the matrix of equation (1)
 - 2: **for** $c = 1, \dots, q - 1$ **do**
 - 3: **Initialize** the messages $v^{(0)} = (v_{i \rightarrow j}^{(0)})_{(ij) \in E}$
 - 4: **for all** $(i \rightarrow j) \in \vec{E}$ **do**
 - 5: **if** $i \in \mathcal{L}$ and $\sigma_i = c$ **then** $v_{i \rightarrow j}^{(0)} \leftarrow 1$
 - 6: **else if** $i \in \mathcal{L}$ and $\sigma_i \neq c$ **then** $v_{i \rightarrow j}^{(0)} \leftarrow -1$
 - 7: **else** $v_{i \rightarrow j}^{(0)} \leftarrow \pm 1$ uniformly at random
 - 8: **Iterate** for $k = 1, \dots, k_{\max}$
 - 9: $v^{(k)} \leftarrow B_c v^{(k-1)}$
 - 10: **Pool** the messages in a vector $\hat{v}_c \in \mathbb{R}^n$ with
 - 11: entries $(\hat{v}_{i,c})_{i \in [n]}$
 - 12: **for** $i \in [n]$ **do** $\hat{v}_{i,c} \leftarrow \sum_{l \in \partial i} w_{il} v_{l \rightarrow i}^{(k_{\max})}$
 - 13: **Deflate** B_c
 - 14: $B_{c+1} \leftarrow B_c - \frac{B_c v^{(k_{\max})} v^{(k_{\max})\top} B_c}{v^{(k_{\max})\top} B_c v^{(k_{\max})}}$
 - 15: **Concatenate** $\hat{V} \leftarrow [\hat{v}_1 | \dots | \hat{v}_{q-1}] \in \mathbb{R}^{n \times (q-1)}$
 - 16: **Output** the assignments $(\hat{\sigma}_i)_{i \in [n]} \leftarrow \text{kmeans}(\hat{V})$
-

3. Numerical simulations

In addition to the theoretical guarantees presented in the previous section, we ran numerical simulations on two toy datasets consisting of 2 and 4 Gaussian blobs (Figure 1), and two subsets of the MNIST dataset [23] consisting respectively of the digits in $\{0, 1\}$ and $\{0, 1, 2\}$ (Figure 2). We also considered the 20 Newsgroups text dataset [24], consisting of text documents organized in 20 topics, of which we selected 2 for our experiments of Figure 3. All three examples differ considerably from the model we have studied analytically. In particular, the random similarities are not identically distributed conditioned on the true labels of the items, but depend on latent variables, such as the distance to the center of the Gaussian, in the case of Figure 1. Additionally, in the case of the MNIST dataset of Figure 2, the clusters have different sizes (e.g. 6903 for the 0's and 7877 for the 1's). Nevertheless, we find that our algorithm performs well, and outperforms the popular label propagation algorithm [8] in a wide range of values of the sampling rate α .

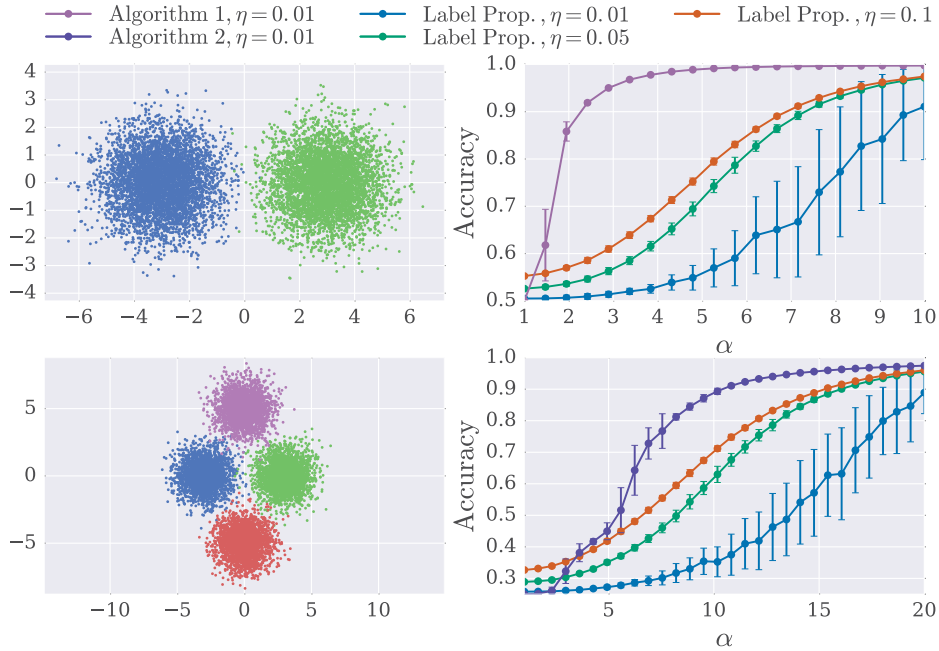


Figure 1. Performance of algorithms 1 and 2 compared to label propagation on a toy dataset in two dimensions. The left panel shows the data, composed of $n = 10^4$ points, with their true labels. The right panel shows the clustering performance on random subsamples of the complete similarity graph. Each point is averaged over 100 realizations. The accuracy is defined as the fraction of correctly labeled points. We set the maximum number of iterations of our algorithms to $k_{\max} = 30$. α is the average degree of the Erdős-Rényi random graph G , and η is the fraction of labeled data. For all methods, we used the same similarity function $s_{ij} = \exp -d_{ij}^2/\sigma^2$ where d_{ij} is the Euclidean distance between points i and j and σ^2 is a scaling factor which we set to the empirical mean of the observed squared distances d_{ij}^2 . For algorithms 1 and 2, we used the weighting function $w(s) := s - \bar{s}$ (i.e. we simply center the similarities, see text). Label propagation is run on the random similarity graph G . We note that Algorithm 1 has the same performance as label propagation with ten times fewer labeled data and with only a number of similarities of the order of three times the number of points.

In all cases, we find that the accuracy achieved by Algorithms 1 and 2 is an increasing function of α , rapidly reaching a plateau at a limiting accuracy. This limiting accuracy is also reached by label propagation and depends on the choice of the similarity function. But the most important observation is the rate of convergence of the accuracy to this limiting value, as a function of α . Indeed, on these simple datasets, it is enough to compute, for each item, their similarity with a few randomly chosen other items to reach an accuracy within a few percents of the limiting accuracy allowed by the quality of the similarity function. As a consequence, similarity-based clustering can be significantly sped up. For example, we note that the semi-supervised clustering of the 0's and 1's of the MNIST dataset (representing $n = 14780$ points in dimension 784), from 1% labeled data, and to an accuracy greater than 96% requires $\alpha \approx 6$ (see figure 2), and runs on a laptop in 2 seconds, including the computation of the randomly chosen similarities. Additionally, in contrast with our algorithms, we find that in the strongly undersampled regime (small α), the performance of label propagation depends strongly on the fraction η of available labeled data. We find in particular that Algorithms 1 and 2 outperform label propagation even starting from ten times fewer labeled data.

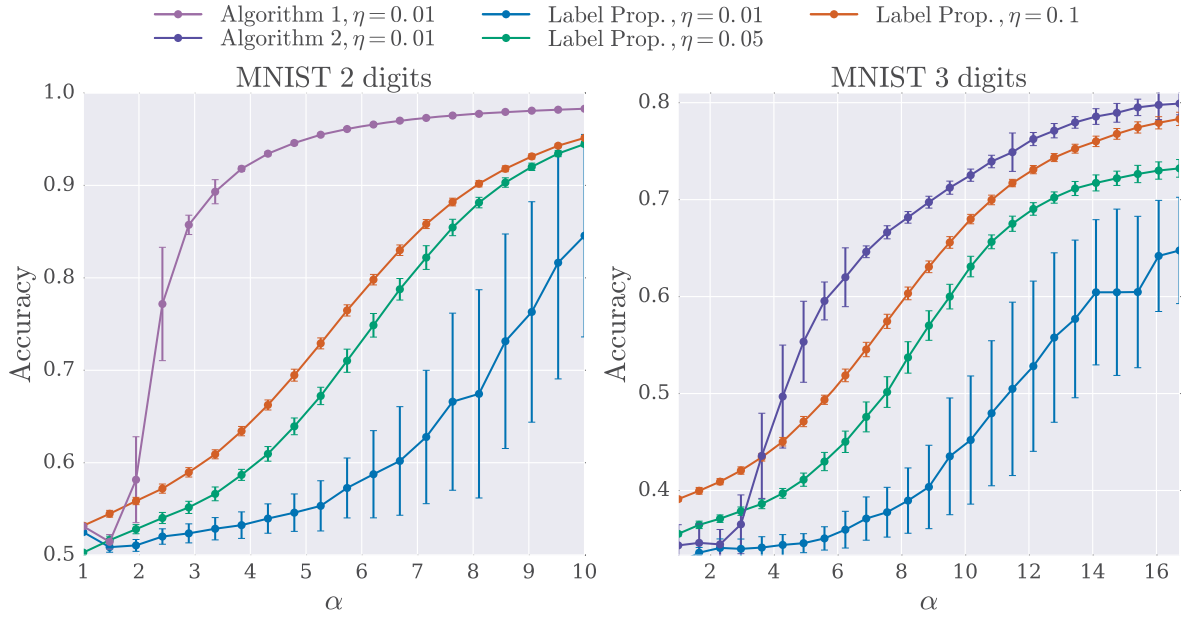


Figure 2. Performance of algorithms 1 and 2 compared to label propagation on a subset of the MNIST dataset. The left panel corresponds the set of 0’s and 1’s ($n = 14780$ samples) while the right panel corresponds the 0’s,1’s and 2’s ($n = 21770$). All the parameters are the same as in the caption of figure 1, except that we used the Cosine distance in place of the Euclidean one.

4. Proof of Theorem 1

Consider the model introduced in Section 2 for the case of two clusters. We will bound the probability of error on a randomly chosen node, and (3) will follow. Denote by I an integer drawn uniformly at random from $[n]$, and by $\hat{\sigma}_I^{(k)}$ the decision variable after k iterations of Algorithm 1. We are interested in the probability of error at node I conditioned on the true label of node I . As noted previously, the algorithm is local in the sense that $\hat{\sigma}_I^{(k)}$ depends only on the messages in the neighborhood of I consisting of all the nodes and edges of G that are at most $k + 1$ steps away from I . Thanks to the total variation distance bound between the law of this neighborhood and a weighted Galton-Watson branching process (see prop. 31 in [20]), we can use a standard coupling argument to reduce the analysis of our algorithm on G to its analysis on a branching process and this will introduce the error term $C \frac{\alpha^{k+1} \log n}{\sqrt{n}}$ in (3). Using the recursive structure of the branching process allows us to determine the law of $\hat{v}_I = \hat{v}_\sigma^{(k)}$ conditioned on the type of node I being $\sigma = \pm 1$,

$$\hat{v}_\sigma^{(k)} \stackrel{\mathcal{D}}{=} \sum_{i=1}^{d_1} w_{i,\text{in}} v_{i,\sigma}^{(k)} + \sum_{i=1}^{d_2} w_{i,\text{out}} v_{i,-\sigma}^{(k)}, \quad (9)$$

where $\stackrel{\mathcal{D}}{=}$ denotes equality in distribution. The random variables $v_{i,\sigma}^{(k)}$ for $\sigma = \pm 1$ have the same distribution as the message $v_{i \rightarrow j}^{(k)}$ after k iterations of the algorithm, for a randomly chosen edge ($i \rightarrow j$), conditioned on the type of node i being σ . They are i.i.d. copies of a random variable

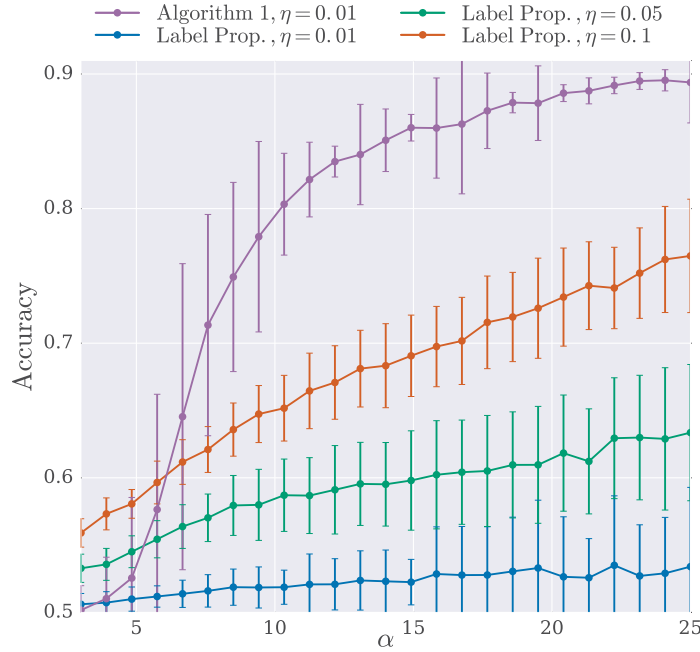


Figure 3. Performance of Algorithm 1 compared to label propagation on a subset of the 20 Newsgroups dataset with $q = 2$ clusters. We consider the topics “misc.forsale” (975 documents) and “soc.religion.christian” (997 documents), which are relatively easy to distinguish, to illustrate the efficiency of our subsampling approach. The resulting dataset consists of $n = 1972$ text documents, to which we applied a standard “tf-idf” transformation (after stemming, and using word unigrams) to obtain a vector representation of the documents. We used the same similarity (based on the Cosine distance) and weighting functions as in figure 2.

$v_{\sigma}^{(k)}$ whose distribution is defined recursively, for $l \geq 0$ and $\sigma = \pm 1$, through

$$v_{\sigma}^{(l+1)} \stackrel{D}{=} \sum_{i=1}^{d_1} w_{i,\text{in}} v_{i,\sigma}^{(l)} + \sum_{i=1}^{d_2} w_{i,\text{out}} v_{i,-\sigma}^{(l)}. \quad (10)$$

In equations (9), d_1 and d_2 are two independent random variables with a Poisson distribution of mean $\alpha/2$, and $w_{i,\text{in}}$ (resp. $w_{i,\text{out}}$) are i.i.d copies of w_{in} (resp. w_{out}) whose distribution is the same as the weights w_{ij} , conditioned on $\sigma_i = \sigma_j$ (resp. $\sigma_i \neq \sigma_j$). Note in particular that $\hat{v}_{\sigma}^{(k)}$ has the same distribution as $v_{\sigma}^{(k+1)}$.

Theorem 1 will follow by analyzing the evolution of the first and second moments of the distribution of $v_{+1}^{(k+1)}$ and $v_{-1}^{(k+1)}$. Equations (10) can be used to derive recursive formulas for the first and second moments. In particular, the expected values verify the following linear system

$$\begin{pmatrix} \mathbb{E}[v_{+1}^{(l+1)}] \\ \mathbb{E}[v_{-1}^{(l+1)}] \end{pmatrix} = \frac{\alpha}{2} \begin{pmatrix} \mathbb{E}[w_{\text{in}}] & \mathbb{E}[w_{\text{out}}] \\ \mathbb{E}[w_{\text{out}}] & \mathbb{E}[w_{\text{in}}] \end{pmatrix} \begin{pmatrix} \mathbb{E}[v_{+1}^{(l)}] \\ \mathbb{E}[v_{-1}^{(l)}] \end{pmatrix}. \quad (11)$$

The eigenvalues of this matrix are $\mathbb{E}[w_{\text{in}}] + \mathbb{E}[w_{\text{out}}]$ with eigenvector $(1, 1)^{\top}$, and $\mathbb{E}[w_{\text{in}}] - \mathbb{E}[w_{\text{out}}]$ with eigenvector $(1, -1)^{\top}$. With the assumption of our model, we have $\mathbb{E}[v_{+1}^{(0)}] = \eta = -\mathbb{E}[v_{-1}^{(0)}]$

which is proportional to the second eigenvector. Recalling the definition of $\Delta(w)$ from section 2, we therefore have, for any $l \geq 0$,

$$\mathbb{E} \left[v_{+1}^{(l+1)} \right] = \alpha \Delta(w) \mathbb{E} \left[v_{+1}^{(l)} \right], \quad (12)$$

and $\mathbb{E} \left[v_{-1}^{(l)} \right] = -\mathbb{E} \left[v_{+1}^{(l)} \right]$. With the additional observation that $\mathbb{E} \left[\left(v_{+1}^{(0)} \right)^2 \right] = \mathbb{E} \left[\left(v_{-1}^{(0)} \right)^2 \right] = 1$, a simple induction shows that for any $l \geq 0$, $\mathbb{E} \left[\left(v_{+1}^{(l)} \right)^2 \right] = \mathbb{E} \left[\left(v_{-1}^{(l)} \right)^2 \right]$. Recalling the definition of $\Sigma(w)^2$ from Section 2, we have the recursion

$$\begin{aligned} \mathbb{E} \left[\left(v_{+1}^{(l+1)} \right)^2 \right] &= \alpha^2 \Delta(w)^2 \mathbb{E} \left[v_{+1}^{(l)} \right]^2 \\ &+ \alpha \Sigma(w)^2 \mathbb{E} \left[\left(v_{+1}^{(l)} \right)^2 \right]. \end{aligned} \quad (13)$$

Noting that since $\Delta(w) > 0$, we have $\sigma \mathbb{E} \left[v_{\sigma}^{(k+1)} \right] > 0$ for $\sigma = \pm 1$, the proof of Theorem 1 is concluded by invoking Cantelli's inequality

$$\mathbb{P} \left(\sigma v_{\sigma}^{(k+1)} \leq 0 \right) \leq 1 - r_{k+1}, \quad (14)$$

with, for $l \geq 0$,

$$r_l := \mathbb{E} \left[v_{\sigma}^{(l)} \right]^2 \mathbb{E} \left[\left(v_{\sigma}^{(l)} \right)^2 \right]^{-1}, \quad (15)$$

where r_l is independent of σ , and is shown to verify the recursion (4) by combining (12) and (13).

5. Proof of theorem 2

The proof is adapted from a technique developed by [9]. We show that the random variables $v_{\sigma}^{(l)}$ are sub-exponential by induction on l . A random variable X is said to be sub-exponential if there exist constants $K > 0, a, b$ such that for $|\lambda| < K$

$$\mathbb{E} \left[e^{\lambda X} \right] \leq \exp \left(\lambda a + \lambda^2 b \right). \quad (16)$$

Define $f_{\sigma}^{(l)}(\lambda) := \mathbb{E} \left[e^{\lambda v_{\sigma}^{(l)}} \right]$ for $l \geq 0$ and $\sigma = \pm 1$. We introduce two sequences $(a_l)_{l \geq 0}, (b_l)_{l \geq 0}$ defined recursively by $a_0 = \eta, b_0 = 1/2$ and for $l \geq 0$

$$\begin{aligned} a_{l+1} &= \alpha \Delta(w) a_l, \\ b_{l+1} &= \alpha \Sigma(w)^2 \left(b_l + \frac{3}{2} \max(a_l^2, b_l) \right). \end{aligned} \quad (17)$$

Note that since we assume that $\alpha \Delta(w) > 1$ and $\alpha \Sigma(w)^2 > 1$, both sequences are positive and increasing. In the following, we show that

$$f_{\sigma}^{(k+1)}(\lambda) \leq \exp \left(\sigma \lambda a_{k+1} + \lambda^2 b_{k+1} \right), \quad (18)$$

for $|\lambda| \leq \left(2 \max(a_k, \sqrt{b_k}) \right)^{-1}$. Theorem 2 will follow from the Chernoff bound applied at

$$\lambda_{\sigma}^* = -\sigma \frac{a_{k+1}}{2b_{k+1}} \min \left(1, \frac{\Sigma(w)^2}{\Delta(w)} \right). \quad (19)$$

The fact that $|\lambda_\sigma^*| \leq (2 \max(a_k, \sqrt{b_k}))^{-1}$ follows from (17). Noting that $\sigma \lambda_\sigma^* < 0$ for $\sigma = \pm 1$, the Chernoff bound allows to show

$$\begin{aligned} \mathbb{P}\left(\sigma v_\sigma^{(k+1)} \leq 0\right) &\leq f_\sigma^{(k+1)}(\lambda_\sigma^*) \\ &\leq \exp\left[-\frac{q_{k+1}}{4} \min\left(1, \frac{\Sigma(w)^2}{\Delta(w)}\right)\right], \end{aligned} \quad (20)$$

where $q_{k+1} := a_{k+1}^2/b_{k+1}$ is shown using (17) to verify the recursion (8). We are left to show that $f_\sigma^{(k+1)}(\lambda)$ verifies (18). First, with the choice of initialization in Algorithm 1, we have for any $\lambda \in \mathbb{R}$

$$\begin{aligned} f_{+1}^{(0)}(\lambda) = f_{-1}^{(0)}(-\lambda) &= \frac{1+\eta}{2} \exp(\lambda) + \frac{1-\eta}{2} \exp(-\lambda) \\ &\leq \exp(\eta\lambda + \lambda^2/2), \end{aligned} \quad (21)$$

where we have used the inequality for $x \in [0, 1], \lambda \in \mathbb{R}$

$$xe^\lambda + (1-x)e^{-\lambda} \leq \exp((2x-1)\lambda + \lambda^2/2). \quad (22)$$

Therefore we have shown $f_\sigma^{(0)}(\lambda) \leq \exp(\sigma\lambda a_0 + \lambda^2 b_0)$. Next, let us assume that for some $l \geq 0$ and for any λ such that $|\lambda| \leq (2 \max(a_{l-1}, \sqrt{b_{l-1}}))^{-1}$,

$$f_\sigma^{(l)}(\lambda) \leq \exp(\sigma\lambda a_l + \lambda^2 b_l), \quad (23)$$

with the convention $a_{-1} = b_{-1} = 0$ so that the previous statement is true for any $\lambda \in \mathbb{R}$ if $l = 0$. The density evolution equations (10) imply the following recursion on the moment-generating functions, for any $\lambda \in \mathbb{R}, \sigma = \pm 1$,

$$\begin{aligned} f_\sigma^{(l+1)}(\lambda) &= \exp\left(-\alpha + \frac{\alpha}{2} \left(\mathbb{E}_{w_{\text{in}}} \left[f_\sigma^{(l)}(\lambda w_{\text{in}})\right] \right. \right. \\ &\quad \left. \left. + \mathbb{E}_{w_{\text{out}}} \left[f_{-\sigma}^{(l)}(\lambda w_{\text{out}})\right]\right)\right). \end{aligned} \quad (24)$$

We claim that for $|\lambda| \leq (2 \max(a_l, \sqrt{b_l}))^{-1}$ and for $\sigma = \pm 1$,

$$\begin{aligned} &\frac{1}{2} \left(\mathbb{E}_{w_{\text{in}}} \left[f_\sigma^{(l)}(\lambda w_{\text{in}})\right] + \mathbb{E}_{w_{\text{out}}} \left[f_{-\sigma}^{(l)}(\lambda w_{\text{out}})\right]\right) \\ &\leq 1 + \sigma a_l \Delta(w) \lambda + \lambda^2 \Sigma(w)^2 \left(b_l + \frac{3}{2} \max(a_l^2, b_l)\right) \end{aligned} \quad (25)$$

Injecting equation (25) in the recursion (24) yields $f_\sigma^{(l+1)}(\lambda) \leq \exp(\sigma\lambda a_{l+1} + \lambda^2 b_{l+1})$, for any λ such that $|\lambda| \leq (2 \max(a_l, \sqrt{b_l}))^{-1}$, with a_{l+1}, b_{l+1} defined by (17). The proof is then concluded by induction on $0 \leq l \leq k$. To show (25), we start from the following inequality: for $|a| \leq 3/4$,

$$\exp(a) \leq 1 + a + (2/3)a^2. \quad (26)$$

With $|w| \leq 1$ as per the assumption of theorem 2, for $|\lambda| \leq (2 \max(a_l, \sqrt{b_l}))^{-1}$, we have for $\sigma = \pm 1$ that $|\sigma\lambda w a_l + \lambda^2 w^2 b_l| \leq 3/4$. Additionally, since a_l and b_l are non-decreasing in l , we also have that $|\lambda| \leq (2 \max(a_{l-1}, \sqrt{b_{l-1}}))^{-1}$, so that by our induction hypothesis, for $\sigma = \pm 1$

$$\begin{aligned} f_\sigma^{(l)}(\lambda w) &\leq \exp(\sigma\lambda w a_l + \lambda^2 w^2 b_l) \\ &\leq 1 + \sigma\lambda w a_l + \lambda^2 w^2 b_l + \frac{2}{3} (\sigma\lambda w a_l + \lambda^2 w^2 b_l)^2 \\ &\leq 1 + \sigma\lambda w a_l + \lambda^2 w^2 b_l + \frac{2}{3} \lambda^2 w^2 (a_l + |\lambda| b_l)^2 \\ &\leq 1 + \sigma\lambda w a_l + \lambda^2 w^2 b_l + \frac{3}{2} \lambda^2 w^2 \max(a_l^2, b_l), \end{aligned}$$

where we have used in the last inequality that $(a_l + |\lambda|b_l)^2 \leq 9/4 \max(a_l^2, b_l)$. The claim (25) follows by taking expectation, and the proof is completed.

Acknowledgments

This work has been supported by the ERC under the European Union’s FP7 Grant Agreement 307087-SPARCS and by the French Agence Nationale de la Recherche under reference ANR-11-JS02-005-01 (GAP project).

References

- [1] Ulrike Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395, 2007.
- [2] Petros Drineas, Alan M Frieze, Ravi Kannan, Santosh Vempala, and V Vinay. Clustering in large graphs and matrices. In *SODA*, volume 99, pages 291–299. Citeseer, 1999.
- [3] Dimitris Achlioptas and Frank McSherry. Fast computation of low-rank matrix approximations. *Journal of the ACM (JACM)*, 54(2):9, 2007.
- [4] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [5] A. Saade, M. Lelarge, F. Krzakala, and L. Zdeborov. Clustering from sparse pairwise measurements. pages 780–784, July 2016.
- [6] F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová, and P. Zhang. Spectral redemption in clustering sparse networks. *Proc. Natl. Acad. Sci.*, 110(52):20935–20940, 2013.
- [7] Sugato Basu, Arindam Banerjee, and Raymond Mooney. Semi-supervised clustering by seeding. In *In Proceedings of 19th International Conference on Machine Learning (ICML-2002)*. Citeseer, 2002.
- [8] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, Citeseer, 2002.
- [9] David R Karger, Sewoong Oh, and Devavrat Shah. Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems*, pages 1953–1961, 2011.
- [10] David R Karger, Sewoong Oh, and Devavrat Shah. Efficient crowdsourcing for multi-class labeling. In *ACM SIGMETRICS Performance Evaluation Review*, volume 41, pages 81–92. ACM, 2013.
- [11] Emmanuel Abbe, Afonso S Bandeira, Annina Bracher, and Amit Singer. Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery. *arXiv:1404.4749*, 2014.
- [12] Se-Young Yun and Alexandre Proutiere. Optimal cluster recovery in the labeled stochastic block model. *ArXiv e-prints*, 5, 2015.
- [13] Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *arXiv preprint arXiv:1502.07738*, 2015.
- [14] M. Lelarge, L. Massoulié, and Jiaming Xu. Reconstruction in the labeled stochastic block model. In *Information Theory Workshop (ITW), 2013 IEEE*, pages 1–5, Sept 2013.
- [15] Simon Heimlicher, Marc Lelarge, and Laurent Massoulié. Community detection in the labelled stochastic block model. 09 2012.
- [16] Alaa Saade, Florent Krzakala, Marc Lelarge, and Lenka Zdeborová. Spectral detection in the censored block model. *IEEE International Symposium on Information Theory (ISIT2015)*, to appear, 2015.
- [17] Pan Zhang. Robust spectral detection of global structures in the data by learning a regularization. *arXiv preprint arXiv:1609.02906*, 2016.
- [18] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 369–376. ACM, 2014.
- [19] Varun Kanade, Elchanan Mossel, and Tselil Schramm. Global and local information in clustering labeled block models. 2014.
- [20] Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs. *arXiv:1501.06087*, 2015.
- [21] T Tony Cai, Tengyuan Liang, and Alexander Rakhlin. Inference via message passing on partially labeled stochastic block models. *arXiv preprint arXiv:1603.06923*, 2016.
- [22] Nguyen Duc Thang, Young-Koo Lee, Sungyoung Lee, et al. Deflation-based power iteration clustering. *Applied intelligence*, 39(2):367–385, 2013.
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [24] Ken Lang. Newsweeder: Learning to filter netnews. In *Proceedings of the 12th international conference on machine learning*, pages 331–339, 1995.