Block Model Guided Unsupervised Feature Selection

Zilong Bai

University of California, Davis
zlbai@ucdavis.edu

Hoa Nguyen

University of California, Davis
hoanguyen@ucdavis.edu

Ian Davidson

University of California, Davis
davidson@cs.ucdavis.edu

ABSTRACT

Feature selection is a core area of data mining with a recent innovation of graph-driven unsupervised feature selection for linked data. In this setting we have a dataset Y consisting of n instances each with m features and a corresponding n node graph (whose adjacency matrix is A) with an edge indicating that the two instances are similar. Existing efforts for unsupervised feature selection on attributed networks have explored either directly regenerating the links by solving for f such that $f(y_i, y_j) \approx A_{i,j}$ or finding community structure in A and using the features in Y to predict these communities. However, graph-driven unsupervised feature selection remains an understudied area with respect to exploring more complex guidance. Here we take the novel approach of first building a block model on the graph and then using the block model for feature selection. That is, we discover $FMF^T \approx A$ and then find a subset of features S that induces another graph to preserve both F and M . We call our approach Block Model Guided Unsupervised Feature Selection (BMGUFS). Experimental results show that our method outperforms the state of the art on several real-world public datasets in finding high-quality features for clustering.

CCS CONCEPTS

• Computing methodologies → Feature selection.

KEYWORDS

Unsupervised Feature Selection; Attributed Networks; Block Model

ACM Reference Format:

Zilong Bai, Hoa Nguyen, and Ian Davidson. 2020. Block Model Guided Unsupervised Feature Selection. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20)*, August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3394486.3403173>

1 INTRODUCTION

The area of feature selection is a critical initial step in data mining and vital for its success. It has been extensively studied [21] with a recent innovation of graph driven feature selection where in addition to an m featured data set Y of n instances, we are given an n node graph whose adjacency matrix between instances is A . Here the graph represents instance similarity such that if $A_{a,b} \geq A_{i,j}$

then instances a and b are more similar than instances i and j . This allows a rich source of guidance for the feature selection process.

Such a setting is not unusual in modern data mining particularly given the proliferation of attributed networks in various domains ranging from social media (e.g., Twitter [39]) to biochemistry (e.g., protein-protein interacting networks [32]). In these settings the nodes are accompanied by a collection of features (an $n \times m$ feature matrix Y) in addition to relational network topology (an $n \times n$ adjacency matrix A). A challenge in these domains is that the nodal attributes can be a noisy/irrelevant or even redundant high-dimensional feature space. This can yield suboptimal solutions if we assume all the features associated with the nodes and the graph structure are complementary [33, 46].

Existing work to address this challenge takes two broad directions to make use of the graph. The first (micro-level) is to learn a function that maps the feature vectors of two instances to a value that approximates their edge weight in the graph, that is $f(y_i, y_j) \approx A_{i,j}$ (e.g., [22, 40, 41]). A second direction (macro-level) includes finding communities from A either explicitly (e.g. [36]) or implicitly (e.g., [23]) and selecting features to predict them. Instead, we take the novel approach of finding a block model $FMF^T \approx A$ and use the block model F and M to guide the feature selection. This is different from existing work in two ways. Firstly, clustering and block modeling are not the same, as in block modeling two instances are placed in the same block if they are structurally equivalent (e.g., second-order proximity [37, 43]), not if they belong to the same densely connected subgraph (i.e., intra-community proximity [12, 43]). Secondly, a block model effectively denoises the graph and hence removes noisy edges. See Figure 1 for an illustration of our work.

Our major contributions are:

- (1) We propose a novel block-model driven formulation for feature selection (section 3).
- (2) We derive an effective numerical optimization framework for our formulation (section 4).
- (3) We empirically demonstrate the usefulness of our method and investigate its potential via extensive experiments on several real-world public datasets (section 5).
 - (a) We demonstrate the effectiveness of our method in finding high-quality features to facilitate K-means clustering. Our method outperforms the baselines on various real-world public datasets (section 5.3).
 - (b) We conduct in-depth analysis on the sensitivity of our method w.r.t. the block models generated from the structural graph to gain insights for future endeavor beyond our explorations (section 5.4).

We begin the rest of the paper by presenting the problem setting in section 2. We then formulate our BMGUFS as an optimization problem in section 3 and derive a highly effective algorithm in section 4. We present the results of our extensive experiments in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7998-4/20/08...\$15.00

<https://doi.org/10.1145/3394486.3403173>

Figure 1: Diagram of our proposed BMGUFs. In the above attributed network $G(V, E, Y)$, we first compute the block model F and M for the structural graph A . We then select features S to induce graph $\hat{A} = Y \text{diag}(\mathbf{r}) Y^T$ to maximally preserve F and M . Particularly, features in orange (i.e., 2, 4, 5) are selected such that on their induced $\hat{A}$: (1) the given block allocation F minimally violates structural equivalence and (2) the image matrix $\hat{M}$ corresponded to F maximally regenerates M .

section 5. We then briefly review highly related work in section 6. We conclude the paper and discuss future directions in section 7.

2 PROBLEM SETTING

In this section we first establish our notation in Table 1. We then present several concepts as preliminaries of our Block Model Guided Unsupervised Feature Selection problem. We formally define the novel feature selection problem we explore in Problem 1.

We denote matrices as boldface capital letters (e.g., X), vectors as boldface lowercase letters (e.g., $\mathbf{x}$), scalars as regular lowercase letters (e.g., x). We index the i -th entry of vector $\mathbf{x}$ with x_i , the i -th row of X with $X_{i,*}$, the j -th column of this matrix with $X_{*,j}$, the (i, j) entry in X with $X_{i,j}$. We use either X' or X^T to denote the

transpose of X . We use $\text{tr}(X)$ to denote the trace of square matrix X . We follow MATLAB syntax to use $\text{diag}(\bullet)$ for *either* diagonalization when $\bullet$ is a vector $\mathbf{x}$ or extracting the m diagonal entries as a vector when $\bullet$ is a square matrix $X^{m \times m}$. We use $\mathbf{1}$ to denote a vector with all elements being 1, $\bar{\mathbf{1}} = \frac{1}{\|\mathbf{1}\|_2}$. For matrix/vector computations, we use $\otimes$ for Kronecker product, $\odot$ for Hadamard (element-wise) product, and $\oslash$ for element-wise division. Horizontal concatenation of two matrices refers to regular matrix product. We use $X(\mathbf{r})$ to denote a matrix X is a matrix function of $\mathbf{r}$. Function $\text{nnz}(\mathbf{r})$ counts the number of non-zero entries in $\mathbf{r}$.

Definition 2.1 (Attributed Network). An attributed network $G(V, E, Y)$ consists of the set of nodes V , the set of links $E \subset V \times V$, and $Y^{n \times m}$ where $Y_{i,*}$ is the m -dimensional feature/attribute vector of node v_i . The adjacency matrix of the **Structural Graph** (i.e., the raw network topology) is A .

Definition 2.2 (Graph Induced by S). Let S be the subset of d features selected from the original m -dimensional feature space $\mathcal{Y}$. The graph induced by S is defined by the similarity between nodes in S . Formally, let $\mathbf{r} \in \{0, 1\}^m$ indicating the l -th feature is selected iff $r_l = 1$, otherwise 0, the adjacency matrix of the graph induced by S is defined as $\hat{A} = Y \text{diag}(\mathbf{r}) Y^T$ for this paper.

Definition 2.3 (Block Model of Graph A). A block modeling result of graph $A^{n \times n}$ that partitions its node set V into k blocks consists of a block allocation matrix $F \in \{0, 1\}^{n \times k}$ and an image/mixing matrix $M \in \mathbb{R}_+^{k \times k}$, s.t. F, M (approximately) minimize $\|A - FMF^T\|_F$. The image matrix M corresponded to F on graph A is $\arg \min_X \|A - FXF^T\|_F$.

PROBLEM 1 (BLOCK MODEL GUIDED UNSUPERVISED FEATURE SELECTION). **Input:** Feature matrix Y for n nodes in the original feature space $\mathcal{Y}$ with m features, block model F and M precomputed from the adjacency matrix A of the structural graph over the n nodes. **Optimization:** Find a subset of d features S from $\mathcal{Y}$ ($d \ll m$), such that the graph induced from S (i.e., $\hat{A}$) maximally preserves F and M .

Notation	Definition
n	Number of nodes.
m	Number of original features.
d	Number of selected features.
k	Number of blocks in a block model.
$\mathcal{Y}$	Original feature set.
S	Selected feature set.
$Y^{n \times m}$	Feature matrix of all n nodes.
$A^{n \times n}$	Adjacency matrix of the structural graph for the attributed network.
$\hat{A}^{n \times n}$	Adjacency matrix of the similarity graph induced from S .
$F^{n \times k}$	Block allocation matrix with k -blocks stacked in columns.
$M^{k \times k}$	Image matrix on the structural graph.
$\hat{M}^{k \times k}$	Image matrix on the induced graph.
$\mathbf{r} \in \{0, 1\}^m$	The feature selection indicator vector.
$\mathbf{r} \in [0, 1]^m$	Importance scores for all the features.

Table 1: Notations and Definitions

Output: An m -dimensional feature selection indicator vector $\mathbf{r} \in \{0, 1\}^m$ where $r_l = 1$ iff feature l from $\mathcal{Y}$ is in $\mathcal{S}$, 0 otherwise.

3 FORMULATION

In this section we formulate the Block Model Guided Unsupervised Feature Selection as an optimization problem. We aim to find a subset of features $\mathcal{S}$, such that a given block model $\mathbf{F}, \mathbf{M}$ precomputed for the structural graph $\mathbf{A}$ is maximally preserved on the graph $\hat{\mathbf{A}}$ induced by $\mathcal{S}$. This consists of two objectives: (1) block allocation $\mathbf{F}$ minimally violates structural equivalence on $\hat{\mathbf{A}}$ and (2) the image matrix $\hat{\mathbf{M}}$ corresponded to $\mathbf{F}$ on $\hat{\mathbf{A}}$ regenerates the given $\mathbf{M}$ up to scaling. We model the two objectives as $\mathcal{L}_b$ (section 3.1) and $\mathcal{L}_m$ (section 3.2) respectively.

Based on theorem 3.1, $\hat{\mathbf{M}}(\mathbf{r})$ is a matrix function of feature selection vector $\mathbf{r}$ given $\mathbf{Y}$ and $\mathbf{F}$ (equation 1). Therefore, both $\mathcal{L}_b$ and $\mathcal{L}_m$ are functions of $\mathbf{r}$ without involving an independent variable matrix to model the image matrix $\hat{\mathbf{M}}$ corresponded to $\mathbf{F}$ on the induced graph.

$$\begin{aligned} \hat{\mathbf{M}}(\mathbf{r}) &= \underset{\mathbf{X}}{\operatorname{argmin}} \|\mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T - \mathbf{F} \mathbf{X} \mathbf{F}^T\|_F \\ &= (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \end{aligned} \quad (1)$$

THEOREM 3.1 (LEAST SQUARES OPTIMAL $\mathbf{M}$ IN CLOSED FORM). Given $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{F} \in \{0, 1\}^{n \times k}$. If $\mathbf{D} = \mathbf{F}^T \mathbf{F}$ is a diagonal matrix with positive diagonal elements, then

$$\underset{\mathbf{X}}{\operatorname{argmin}} \|\mathbf{A} - \mathbf{F} \mathbf{X} \mathbf{F}^T\|_F^2 = \mathbf{D}^{-1} \mathbf{F}^T \mathbf{A} \mathbf{F} \mathbf{D}^{-1} \quad (2)$$

PROOF. See Appendix A. $\square$

3.1 Preserving Structural Equivalence with $\mathbf{F}$

Here we aim to find a feature subset $\mathcal{S}$ such that on its induced graph $\hat{\mathbf{A}}$, block allocation $\mathbf{F}$ minimally violates the structural equivalence. According to [27], the reconstruction error $\|\hat{\mathbf{A}} - \hat{\mathbf{M}} \mathbf{F}^T\|$ quantifies the violation of structural equivalence in using $\mathbf{F}$ and $\hat{\mathbf{M}}$ to model $\hat{\mathbf{A}}$. Note that the scale of absolute reconstruction error favors fewer entries in $\mathbf{r}$ to be positive, which can yield trivial solutions (e.g., $\mathbf{r} = \mathbf{0}$) instead of exploring more meaningful block models. Therefore, we model the loss term in equation 3 with the relative reconstruction error for $\hat{\mathbf{A}} = \mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T$. In equation 3, $\mathbf{D} = \mathbf{F}^T \mathbf{F}$.

$$\begin{aligned} \mathcal{L}_b(\mathbf{r}) &= \frac{\|\mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T - \mathbf{F} \hat{\mathbf{M}}(\mathbf{r}) \mathbf{F}^T\|_F^2}{\|\mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T\|_F^2} \\ &\stackrel{(Eq. 1)}{=} \frac{\|\mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T - \mathbf{F} \mathbf{D}^{-1} \mathbf{F}^T \mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T \mathbf{F} \mathbf{D}^{-1} \mathbf{F}^T\|_F^2}{\|\mathbf{Y} \operatorname{diag}(\mathbf{r}) \mathbf{Y}^T\|_F^2} \end{aligned} \quad (3)$$

3.2 Regenerating Image Matrix $\mathbf{M}$

Here we aim to find a feature subset $\mathcal{S}$ such that the image matrix $\hat{\mathbf{M}}$ corresponded to $\mathbf{F}$ on the graph $\hat{\mathbf{A}}$ induced by $\mathcal{S}$ (approximately) regenerates the given $\mathbf{M}$. The underlying premise is that given the same block allocation $\mathbf{F}$, we want the block-level similarity (i.e., $\hat{\mathbf{M}}(\mathbf{r})$) on $\hat{\mathbf{A}}$ to respect the block-level connectivity (i.e., $\mathbf{M}$) on the structural graph $\mathbf{A}$. This translates to $\mathbf{C} \odot \hat{\mathbf{M}} \approx \mathbf{M}$ where $\mathbf{C}$ is a scalar that compensates for the scaling difference between the two image matrices. It is challenging to directly model and solve for $\mathbf{C}$

as it is not only unknown but also dynamic as the scale of $\hat{\mathbf{M}}(\mathbf{r})$ changes with $\mathbf{r}$. Therefore, we define distance between $\hat{\mathbf{M}}(\mathbf{r})$ and $\mathbf{M}$ in equation 4 invariant to scaling. We use $\mathcal{D}_{KL}(\bullet || \bullet)$ to denote the KL-divergence [20] between two discrete probabilistic distributions. In this paper we only consider block modeling result whose image matrix $\mathbf{M}$ does not contain absolutely zero elements as the presence of absolute zero entries in $\mathbf{M}$ can pose additional challenges to understanding the stochastic properties of the block model [1].

$$\mathcal{L}_m(\mathbf{r}) = \sum_{i,j \in [k]} \mathcal{D}_{KL}(\mathbf{q}_i(\mathbf{r}) || \mathbf{p}_i)$$

where $\mathbf{p}_i = \mathbf{P}_{i,*}$, $\mathbf{q}_i(\mathbf{r}) = \mathbf{Q}(\mathbf{r})_{i,*}$,

$$\mathbf{P}_{i,j} = \frac{\mathbf{M}_{i,j}}{\sum_{j' \in [k]} \mathbf{M}_{i,j'}}, \quad \mathbf{Q}(\mathbf{r})_{i,j} = \frac{\hat{\mathbf{M}}(\mathbf{r})_{i,j}}{\sum_{j' \in [k]} \hat{\mathbf{M}}(\mathbf{r})_{i,j'}} \quad (4)$$

A Statistical Interpretation. In the given $\mathbf{M}$, $\mathbf{M}_{i,j} \in [0, 1]$ can be interpreted as the empirical probability of having an edge connecting two nodes between blocks $\mathbf{F}_{*,i}$ and $\mathbf{F}_{*,j}$. This induces the conditional probability given $\mathbf{F}_{*,i}$ to connect with $\mathbf{F}_{*,j}$ as $\mathbf{P}_{i,j} = \frac{\mathbf{M}_{i,j}}{\sum_{j' \in [k]} \mathbf{M}_{i,j'}}$.

Assuming $\hat{\mathbf{M}}(\mathbf{r}) \propto \mathbf{M}$, we can define $\mathbf{Q}(\mathbf{r})_{i,j} = \frac{\hat{\mathbf{M}}(\mathbf{r})_{i,j}}{\sum_{j' \in [k]} \hat{\mathbf{M}}(\mathbf{r})_{i,j'}}$ to model

the conditional probability given $\mathbf{F}_{*,i}$ to connect to $\mathbf{F}_{*,j}$ on the induced graph. Thus equation 4 models the overall KL-divergence between the conditional probabilities of connectivity on the original graph $\mathbf{A}$ and the induced graph $\hat{\mathbf{A}}(\mathbf{r})$ at the block-level.

3.3 A Joint Formulation

We aim to holistically utilize both the block allocation $\mathbf{F}$ and the image matrix $\mathbf{M}$ to regularize the macro-level structure of the graph induced by the selected features. Therefore, we combine $\mathcal{L}_b$ and $\mathcal{L}_m$ into a unified optimization framework in equation 5 with an adaptive weighting factor $\hat{\beta} \geq 0$ ¹.

$$\begin{aligned} \underset{\mathbf{r}}{\operatorname{Minimize}} \quad \mathcal{L} &= \mathcal{L}_b + \hat{\beta} \mathcal{L}_m \\ \text{s.t. } \mathbf{r} &\in \{0, 1\}^m, \quad \mathbf{r}^T \mathbf{1} = d \end{aligned} \quad (5)$$

To side step the potential intractability caused by combinatorial optimization, we relax the domain of $\mathbf{r}$ from $\{0, 1\}^m$ to $[0, 1]^m$. The resulting $\mathbf{r}$ can be interpreted as importance scores for ranking the features. We then follow the convention of [22] to rewrite the cardinality constraint $\mathbf{r}^T \mathbf{1} = d$ in the Lagrangian, resulting in the following constrained optimization problem with l_1 norm regularization (where γ denotes the weight for sparsity penalty). We further notice that both $\mathcal{L}_b$ and $\mathcal{L}_m$ are invariant to the l_2 norm of $\mathbf{r}$. Therefore, we introduce l_2 norm constraint $\|\mathbf{r}\|_2 = 1$ to confine the search domain for our gradient-descent based algorithm. Equation 6 presents the resulting relaxed formulation.

$$\begin{aligned} \underset{\mathbf{r}}{\operatorname{Minimize}} \quad \mathcal{L} &= \mathcal{L}_b + \hat{\beta} \mathcal{L}_m + \gamma \|\mathbf{r}\|_1 \\ \text{s.t. } \mathbf{r} &\geq \mathbf{0}, \quad \|\mathbf{r}\|_2 = 1 \end{aligned} \quad (6)$$

4 SOLVER

In this section we derive an effective solver for equation 6 to find a feature selection vector $\mathbf{r}$ given block model $\mathbf{F}, \mathbf{M}$. Firstly, we

¹This is not the hyper-parameter $\bar{\beta} \in [0, 1]$ for our algorithm 1.

compute the partial derivatives of $\mathcal{L}_b$ and $\mathcal{L}_m$ w.r.t. $\mathbf{r}$. We then suggest an update rule for $\mathbf{r}$ based on a weighted combination of the *normalized* gradients. We summarize our optimization framework in algorithm 1.

The derivation of $\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}}$ is relatively straightforward - we induce equation 8 from equation 7.

$$\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}} = \frac{1}{\|\mathbf{Y} \mathbf{R} \mathbf{Y}'\|_F^4} (\|\mathbf{Y} \mathbf{R} \mathbf{Y}'\|_F^2 \frac{\partial \|\mathbf{Y} \text{diag}(\mathbf{r}) \mathbf{Y}^T - \mathbf{F} \hat{\mathbf{M}}(\mathbf{r}) \mathbf{F}^T\|_F^2}{\partial \mathbf{r}} - \|\mathbf{Y} \text{diag}(\mathbf{r}) \mathbf{Y}^T - \mathbf{F} \hat{\mathbf{M}}(\mathbf{r}) \mathbf{F}^T\|_F^2 \frac{\partial \|\mathbf{Y} \mathbf{R} \mathbf{Y}'\|_F^2}{\partial \mathbf{r}}) \quad (7)$$

Gradient of $\mathcal{L}_b$ over $\mathbf{r}$

$$\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}} = \frac{2 \text{diag}(\mathbf{Y}' \mathbf{Y} \mathbf{R} \mathbf{Y}' \mathbf{Y} + \hat{\mathbf{Y}}' \hat{\mathbf{Y}} \mathbf{R} \hat{\mathbf{Y}}' \hat{\mathbf{Y}} - 2 \hat{\mathbf{Y}}' \mathbf{Y} \mathbf{R} \mathbf{Y}' \hat{\mathbf{Y}})}{\|\mathbf{Y} \mathbf{R} \mathbf{Y}'\|_F^2} - \frac{2 \mathcal{L}_b \text{diag}(\mathbf{Y}' \mathbf{Y} \mathbf{R} \mathbf{Y}' \mathbf{Y})}{\|\mathbf{Y} \mathbf{R} \mathbf{Y}'\|_F^2}$$

where $\mathbf{R} = \text{diag}(\mathbf{r})$, $\hat{\mathbf{Y}} = \mathbf{F} \mathbf{D}^{-1} \mathbf{F}' \mathbf{Y}$, $\mathbf{D} = \mathbf{F}^T \mathbf{F}$

(8)

To derive $\frac{\partial \mathcal{L}_m}{\partial \mathbf{r}}$ we first compute $\frac{\partial \mathcal{L}_m}{\partial \mathbf{Q}}$ (equation 9) and $\frac{\partial \mathbf{Q}_{i,j}}{\partial r_l}$ (equation 10).

$$\frac{\partial \mathcal{L}_m}{\partial \mathbf{Q}} = \log \frac{\mathbf{Q}}{\mathbf{P}} + \mathbf{P} \quad (9)$$

$$\frac{\partial \mathbf{Q}_{i,j}}{\partial r_l} = \frac{1}{\sum_{j'} \hat{\mathbf{M}}_{i,j'}} [\bar{\mathbf{D}}_{i,l} \bar{\mathbf{D}}_{j,l}^T - \mathbf{Q}_{i,j} \bar{\mathbf{D}}_{i,l} \Sigma \bar{\mathbf{D}}_{j,l}^T] \quad (10)$$

where $\hat{\mathbf{M}} = \bar{\mathbf{D}} \text{diag}(\mathbf{r}) \bar{\mathbf{D}}^T$, $\bar{\mathbf{D}} = \mathbf{D}^{-1} \mathbf{F}^T \mathbf{Y}$

Thus, we have the gradient of $\mathcal{L}_m$ over $r_l, \forall l \in [m]$ based on chain rule of partial derivations²:

Gradient of $\mathcal{L}_m$ over r_l

$$\begin{aligned} \frac{\partial \mathcal{L}_m}{\partial r_l} &= \text{tr} \left(\left[\log \frac{\mathbf{Q}}{\mathbf{P}} + \mathbf{P} \right]^T \frac{\partial \mathbf{Q}}{\partial r_l} \right) \\ \frac{\partial \mathbf{Q}}{\partial r_l} &= \text{diag} \left(\frac{1}{\hat{\mathbf{M}} \mathbf{1}} \right) [\bar{\mathbf{D}}_{*,l} \bar{\mathbf{D}}_{*,l}^T - (\bar{\mathbf{D}}_{*,l}^T \mathbf{1}) \text{diag}(\bar{\mathbf{D}}_{*,l}) \mathbf{Q}] \end{aligned} \quad (11)$$

where $\mathbf{Q}$ and $\mathbf{P}$ are computed according to equation 4.

4.1 Combining the Two Gradients to Update $\mathbf{r}$

In this section, we compute the gradient w.r.t. $\mathbf{r}$ to simultaneously optimize $\mathcal{L}_b$ and $\mathcal{L}_m$. It is conventional to combine the two gradients as $\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}} + \beta \frac{\partial \mathcal{L}_m}{\partial \mathbf{r}}$ with a constant hyper-parameter β . However, we observe that objective $\mathcal{L}_b$ can dominate the minimization of $\mathcal{L}_b + \beta \mathcal{L}_m$. This can lead to *increased* $\mathcal{L}_m$ unless β is extremely large. According to our empirical study, the increment of $\mathcal{L}_m$ affects the quality of selected features, and it is difficult to search for a proper $\beta \in \mathbb{R}_+$. We alleviate this issue with a heuristic that

²In practice we add an ignorable positive scalar $\delta = 10^{-6}$ to $\mathbf{M}_{i,j}$, $\hat{\mathbf{M}}_{i,j}$, $\forall i, j \in [k]$ in computing $\mathbf{P}$, $\mathbf{Q}$ and $\frac{\partial \mathbf{Q}}{\partial r_l}$, $\forall l \in [m]$ to avoid numerical instability. The experimental results of this paper are indifferent to δ being 0 or 10^{-6} .

combines the *normalized* gradients proportionally according to a user-specified **composition ratio** $\tilde{\beta} \in [0, 1]$. With more confined hyper-parameter search space, this strategy is simple yet highly effective in practice to control the optimization of $\mathcal{L}_b$ and $\mathcal{L}_m$.

Equation 12 computes the combined gradient with sparsity penalty weight γ . We use Projected Gradient Descent (PGD) followed by rescaling/normalization to update $\mathbf{r}$, such that $\mathbf{r}$ satisfies both the non-negativity and l_2 norm constraints per iteration. The updating is formally defined as equation 13 performed in order, where $\eta^{(t)}$ is the step size at the t -th iteration.

$$\frac{\partial \mathcal{L}}{\partial \mathbf{r}} = \frac{(1 - \tilde{\beta})}{\|\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}}\|_2} \frac{\partial \mathcal{L}_b}{\partial \mathbf{r}} + \frac{\tilde{\beta}}{\|\frac{\partial \mathcal{L}_m}{\partial \mathbf{r}}\|_2} \frac{\partial \mathcal{L}_m}{\partial \mathbf{r}} + \gamma \bar{\mathbf{1}} \quad (12)$$

$$\mathbf{r} = \mathbf{r}^{(t)} - \eta^{(t)} \frac{\partial \mathcal{L}}{\partial \mathbf{r}};$$

$$r_l \leftarrow \text{Max}(r_l, 0), \forall l \in [m]; \quad (13)$$

$$\mathbf{r}^{(t+1)} \leftarrow \frac{\mathbf{r}}{\|\mathbf{r}\|_2}$$

We summarize our optimization framework for equation 6 in algorithm 1. Given d , a specific number of selected features, we select the top d features with the largest importance scores in $\mathbf{r} \in [0, 1]^m$. We empirically demonstrate the convergence of $\mathcal{L}_b + \mathcal{L}_m$ with properly set $\tilde{\beta}$ in section 5.6. Interestingly, we observe that our method can select high-quality features when $\tilde{\beta}$ effectively reduces $\mathcal{L}_m$ (Figure 6a). We leave theoretical investigation on using equations 12 and 13 as general purpose optimization technique to future endeavors.

Algorithm 1 Block Model Guided Unsupervised Feature Selection

Require: Block allocation $\mathbf{F}$, image matrix $\mathbf{M}$ (precomputed from structural graph $\mathbf{A}$), feature matrix $\mathbf{Y}$, ratio for combining gradients $\tilde{\beta}$, sparsity regularization weight γ , maximum iterations maxI , (constant) step size η .

- 1: Initialize feature selection vector $\mathbf{r} = \frac{\mathbf{1}}{\|\mathbf{1}\|_2}$.
- 2: **while** Termination Condition Unsatisfied **do**
- 3: Compute $\frac{\partial \mathcal{L}_b}{\partial \mathbf{r}}$ with equation 8 and $\frac{\partial \mathcal{L}_m}{\partial \mathbf{r}}$ with equation 11.
- 4: Compute gradient $\frac{\partial \mathcal{L}}{\partial \mathbf{r}}$ with equation 12 given $\tilde{\beta}$ and γ .
- 5: Use $\frac{\partial \mathcal{L}}{\partial \mathbf{r}}$ and η to update $\mathbf{r}$ based on PGD and rescaling to satisfy constraints with equation 13.
- 6: **end while**
- 7: **return** : Feature selection vector $\mathbf{r}$.

Computational Complexity Analysis. The computational cost of our algorithm for computing the gradients in each iteration is given by $O(k^3 m + m^3)$ where m is the original number of features, k is the number of blocks in the block model (a very small integer). The number of nodes is irrelevant to the computational cost in each iteration if we precompute constant matrices to avoid redundant computations.

5 EXPERIMENTS

³ In this section, we extensively evaluate our method on various real-world public datasets to address the following questions:

³Source codes available in <https://github.com/ZilongBai/KDD2020BMGUFS> for reproducibility.

- **Q1. Effectiveness of our method (section 5.3).** Can our method find high-quality features to facilitate downstream application (see Figure 2)?
- **Q2. Sensitivity to Block Model Guidance (section 5.4).** The question is multi-facet and we focus on the following two in this paper due to space limitations:
 - (1) Is our method sensitive to perturbations in block model guidance as input to our algorithm 1 (see Figure 3)?
 - (2) Can **different** block models generated from the **same** structural graph offer different guidance (see Figure 4)?
- **Q3. Sensitivity to Parameter Selection (section 5.5).** How do the composition ratio β and sparsity penalty γ influence the clustering performance of the features selected by our method (see Figure 5)?
- **Q4. Solver Inspection (section 5.6).** How does the composition ratio β influence the optimization process of the objective function of our model (see Figure 6)?

5.1 Experimental Settings

Datasets. We test our method on three real-world public datasets: BlogCatalog [17], Citeseer (sparse graph) [18, 34] and Cora [18, 34]. Table 2 summarizes basic statistics of the three datasets. See Appendix B.1 for details on dataset preprocessing.

Baselines. We compare with the following baselines to demonstrate the effectiveness of our method (Q1). We use the source codes provided by the paper authors to reproduce MMPOP and NetFS. We apply methods in `scikit-feature`[21] to obtain the results of LapScore, SPEC, and NDFS. See Appendix B.4 for links to their source codes and settings of their hyper-parameters.

- **All features.**
- **LapScore** [14] evaluates the importance of a feature based on its power of preserving locality.
- **SPEC** [45] proposes a unified framework for feature selection based on spectral graph theory.
- **NDFS** [24] jointly learns cluster labels via spectral clustering and feature selection matrix with $l_{2,1}$ -norm regularization.
- **MMPOP**[41] selects features to maximally preserve local partial order on the structural graph.
- **NetFS** [23]⁴ embeds latent representation learning that respects network clustering into feature selection.

Metrics for Performance Evaluation. We follow the convention [21, 22] to use K-means clustering on selected features (after normalization) as downstream application to evaluate the quality of selected features. We follow the typical settings in [21, 42] to use Accuracy (ACC in equation 19) and Normalized Mutual Information (NMI in equation 20) as performance metrics. See Appendix B.3 for their detailed definitions. Conventionally, the higher ACC and NMI, the higher quality of the features. We report the mean result after 20 runs of K-means to compensate for randomness.

⁴According to the empirical evaluation in [22]: (1) NetFS [23] can achieve state-of-the-art ACC and NMI on BlogCatalog at $d = 200$ (even better than itself at $d \in \{600, 1000\}$). (2) ADAPT [22] and NetFS can achieve similar performance - better than their baseline methods - on various datasets w.r.t. varying number of selected features.

Statistic	BlogCatalog	Citeseer	Cora
Nodes #	5196	3312	2708
Links #	171743	4660	5278
Features #	8189	3703	1433
Classes #	6	6	7

Table 2: Summary on Statistics of Datasets

5.2 Building Block Models for Structural Graph

As we discuss in the related work (section 6), there exist a plethora of approaches for block modeling. We use the multiplicative update rules for the Orthogonal Nonnegative Matrix tri-Factorization (ON-MtF) formulation (equation 14) proposed by seminal work [7] to generate multiple candidate block models. Since $\mathbf{F}, \mathbf{M}$ are not jointly convex in the formulation, we can harvest multiple (i.e., 10) different block models based on random initializations for each dataset. Each block model is computed with 100 iterations⁵. We convert $\mathbf{F} \in [0, 1]^{n \times k}$ to $\mathbf{F} \in \{0, 1\}^{n \times k}$ by setting the largest entry on each row to 1, others to 0. The number of blocks k is set to the number of classes for each dataset. We then compute $\mathbf{M}$ based on equation 2 in theorem 3.1. The 10 different block models are identified by $\#i, i \in [10]$ according to the order they were generated.

$$\underset{\mathbf{F} \geq 0, \mathbf{M} \geq 0}{\text{Minimize}} \|\mathbf{A} - \mathbf{F}\mathbf{M}\mathbf{F}^T\|_F \text{ s.t., } \mathbf{F}^T\mathbf{F} = \mathbf{I} \quad (14)$$

We define relative reconstruction error (RRE) of using block model $\mathbf{F}, \mathbf{M}$ for reconstructing adjacency matrix $\mathbf{A}$ in equation 15 to facilitate block model selection *before* running our algorithm 1.

$$RRE(\mathbf{F}, \mathbf{M}, \mathbf{A}) = \frac{\|\mathbf{A} - \mathbf{F}\mathbf{M}\mathbf{F}^T\|_F}{\|\mathbf{A}\|_F} \quad (15)$$

5.3 Effectiveness of Our Method

We demonstrate the effectiveness of our method by comparing against baseline methods in K-means clustering performance on the selected features. We vary the number of selected features $d \in \{16, 64, 128, 200, 600\}$. The comparison results are in Figure 2. We follow the principles from its original paper to set hyper-parameters for each baseline method. We leave the discussion on model selection to latter sections and report the results of our method with the following parameter setting:

- We fix the composition ratio $\beta = 0.6$ based on observations in sections 5.5 and 5.6.
- We set γ via grid search in $\{0, 0.5, 1, \dots, 5\}$ while $nnz(\mathbf{r}) \geq d$. This is because overly strong γ can force our method to generate too many absolutely zero entries in $\mathbf{r}$ to pick top d features based on non-zero entries in $\mathbf{r}$.
- We report the results of two block models amongst the 10 candidates for each dataset. One is chosen for having the lowest RRE (i.e., *BMGUFs** in Figure 2), thus selected before running our algorithm 1. The other is selected via grid-search (i.e., *BMGUFs^s* in Figure 2).

Figure 2 demonstrates the superiority of our method over baselines in experiments. Specifically, we observe:

⁵Multiplicative update rules are recognized to converge slowly in solving NMF formulations [25]. We set the maximum iterations to 100 where the objective function does not observably decrease.

- Our BMGUFS selects features that achieve better clustering performance in basically all the investigated cases than the baselines. We constantly outperform our major competitor method NetFS in various settings.

We conjecture the superiority of our method on all the investigated datasets with the following explanations:

- (1) Block model of the structural graph provides more robust guidance against noisy links on real-world networks than detailed links and disconnections. Therefore our method outperforms the methods that use micro-level guidance (e.g., MMPOP [41]).
 - (2) Structural equivalence appreciated by block models is more informative than intra-community proximity between nodes to guide feature selection on the investigated datasets. Therefore our method outperforms the methods guided by macro-level graph structure based on community analysis (e.g., NetFS [23]).
- Our method achieves predominant ACC and NMI on each dataset at extremely small number of features. Specifically, our method outperforms the clustering results using all features by over 11% in ACC with only $d = 16$ features on BlogCatalog, whereas other baseline methods fail to surpass the performance of **all features** with such a small amount of features. This highlights the power of our method in both finding high-quality features and dimension reduction.
 - The clustering performance of baseline methods that only consider the feature matrix, i.e., LapScore, SPEC, and NDFS, are consistently suboptimal to ours; however, they can outperform other baselines that incorporate graph structure in some cases. This supports the underlying premise of our work that *block model can be a better way than other approaches to extract guidance from the structural graph for unsupervised feature selection*.

5.4 Sensitivity to Block Model Guidance

In this section we explore the sensitivity of our method w.r.t. the block model guidance from two perspectives: (P1) sensitivity to perturbations in the input block model and (P2) sensitivity to different block models generated from the same structural graph. We present the results on BlogCatalog as similar patterns exist on other datasets. We fix $\gamma = 2$, $\tilde{\beta} = 0.6$ for this section. Figure 3 shows our algorithm 1 is relatively robust against small perturbations in the input block model. Figure 4 shows the feature selection varies with different block models generated from the same graph, yet RRE is a reasonable criteria to select block model from multiple candidates before running our algorithm (see Figure 2).

(P1) We select the block model of lowest RRE as the base block model F , M . We introduce different levels of artificial perturbations (i.e., 5% and 10%) by randomly selecting the given percentage of nodes and modifying their block memberships (i.e., random re-allocation). We explore two situations: (a) only perturb F but keep the original M and (b) perturb F and recompute M for the structural graph A . We measure the difference between the feature selection vectors r generated by the perturbed block models and r_0 generated by the base block model with cosine distance, and summarize the

results as box plots in Figure 3. We observe that our method is robust to small perturbations in block allocation.

(P2) Different block models can be generated from the *same* graph (e.g., from different local optima of equation 14). In Figure 4 (dark blue bars), we compare the K-means clustering performance on features selected by our method guided by the 10 candidate block models generated from the structural graph of BlogCatalog. We observe notable variance in clustering performance on selected features of different block models. Nonetheless, the block model with the lowest RRE (e.g., $ID = 3$ in Figure 4) can guide our feature selection to highly-competitive (or even better) clustering performance in comparison to baselines (see *BMGUFS** in Figure 2). This demonstrates the benefit of generating multiple block models as candidates to guide feature selection. We suggest to choose the block model with the lowest RRE as the guidance of our BMGUFS to alleviate grid-search in scenarios sensitive to computational cost.

We also explore whether (i) the quality of selected features is correlated with (ii) the accuracy of using block model to predict the ground-truth labels. We use the block allocation F as node clustering result to predict the ground-truth labels, and present the results in the yellow bars in Figure 4. We observe no direct correlation between (i) and (ii).

5.5 Sensitivity to Parameter Selection

We investigate the sensitivity of our BMGUFS to the two key hyper-parameters, composition ratio $\tilde{\beta}$ and sparsity penalty γ . We vary $\tilde{\beta} \in \{0, 0.1, 0.2, \dots, 0.9, 1\}$ and $\gamma \in \{0, 0.5, 1, \dots, 4, 4.5, 5\}$ in our experiments. We make two observations: (1) $\tilde{\beta} = 0.6$ is preferred by various datasets and (2) it is practical to perform grid-search to find optimal γ .

In Figure 5, we show the results of the top $d = 16$ features selected by our method on BlogCatalog (first row) and Cora (second row) as representative examples. Citeseer shows similar pattern as BlogCatalog. We set $d = 16$ features to investigate a wider range of γ as we witness stronger γ can yield r with less non-zero entries than the required number of selected features. Nevertheless, our experiments cover cases where γ is too large to ensure $nnz(r) \geq d$. We report the clustering performance as 0's if $nnz(r) < d$.⁶

Figure 5 shows a clear transition pattern around $\tilde{\beta} = 0.5$ for both datasets. The clustering performance sustains at a relatively high level for all investigated $\tilde{\beta} \geq 0.6$ when $nnz(r) \geq d$. Varying γ does not induce significant change in the clustering performance at $\tilde{\beta} \geq 0.6$, unless it is too large to sustain $nnz(r) \geq d$. We acknowledge that different datasets favor different strength of sparsity (e.g., BlogCatalog favors strong sparsity yet Cora prefers none sparsity regulation). Note that the search space of γ for our method is confined and bounded from both ends, because $\gamma \geq 0$ and γ cannot be too large in order to sustain $nnz(r) \geq d$. Therefore, it is practical to use grid search in practice to pursue better performance.

5.6 Solver Inspection

In this section we provide an empirical study on our heuristic solver. We inspect the values per iteration of our two objectives: $\mathcal{L}_b$, relative reconstruction error of block model on induced graph and, $\mathcal{L}_m$,

⁶We avoid using zero entries in r in ranking the features as they are less informative than the non-zero ones.

Figure 2: Comparing K-means clustering performance of features selected by different unsupervised feature selection methods on different datasets. Top row: ACC (%). Bottom row: NMI (in $[0, 1]$). Datasets from left to right: BlogCatalog, Citeseer, Cora. Our method $BMGUFSS^s$ with grid-search (the rightmost bars) achieves superior performance in basically all the cases. Our method $BMGUFSS^*$ that uses the block model with the lowest RRE (the second rightmost bars) can attain highly-competitive results. Our method demonstrates advantage with extremely small amount of features (i.e., $d = 16$).

Figure 3: Sensitivity analysis of BMGUFSS w.r.t. different levels of perturbations in the input block model. Cosine distance measures the difference between $\hat{r}$ learnt from the perturbed block models and r_0 of the block model without perturbation. Left: only perturb block allocation F with random re-allocation, keep the original image matrix M. Right: perturb block allocation F and adjust M based on the structural graph. Our BMGUFSS is robust to small perturbations in F.

the KL-divergence based distance measure between $\hat{M}$ and the given M. As a representative example, we apply our BMGUFSS at $\gamma = 2$ on BlogCatalog guided by the block model with lowest RRE. We vary

Figure 4: Comparing clustering performance on selected features guided by different block models (in blue bars). Clustering performance varies using *different* block models generated from the *same* structural graph. The prediction accuracy using block allocation (in yellow bars) does *not* directly correlate to feature selection quality.

$\hat{\beta} \in \{0, 0.1, \dots, 1\}$ to demonstrate how it balances the decrement of the two objectives. Figure 6 shows that at $\hat{\beta} \geq 0.6$ (solid curves), $\mathcal{L}_m$ monotonically decreases and $\mathcal{L}_b + \mathcal{L}_m$ converges within 200 iterations. Interestingly, these $\hat{\beta}$ correlate with highly-competitive or state-of-the-art clustering performance in our experiments (see

Figure 5: Parameter sensitivity analysis w.r.t. β and γ for K-means clustering on selected features. While $nnz(r) \geq d$, the clustering performance is generally stable w.r.t. γ ; β shows clear cliff effect around $\beta = 0.5$ but remains stable at $\beta \geq 0.6$.

Figure 6: Inspection on the objective function per iteration with varying β . Left: $\mathcal{L}_m$ vs iterations. Right: $\mathcal{L}_b + \mathcal{L}_m$ vs iterations. When $\beta \geq 0.6$, $\mathcal{L}_m$ and $\mathcal{L}_b + \mathcal{L}_m$ both decrease smoothly and monotonically, and converge within 200 iterations. Compared to Figure 5, this suggests the importance of ensuring the minimization of $\mathcal{L}_m$ for our method to acquire high-quality features for clustering.

previous sections). This suggests the relative importance of optimizing $\mathcal{L}_m$ for our method to acquire high-quality features for clustering.

6 RELATED WORK

In this section we briefly review related work on unsupervised feature selection and discuss applications of block model.

Unsupervised Feature Selection. There have been a lot of efforts for improving unsupervised feature selection performance. Many

different methods have been proposed for solving different problems, e.g. adding $l_{2,1}$ -norm minimization constraint to reduce the redundant or even noisy features [24, 42], algorithms for improving the robustness of graph embedding and sparse spectral regression [35], various methods for adaptive structure learning in different scenarios [8, 22], etc. Concerning different selection strategies, feature selection methods can be broadly categorized [21] as wrapper (e.g., [10, 19]), filter (e.g., [14, 24, 45]), and embedded methods (e.g., [6, 15, 23, 31]).

The concept of graph has long been explored in unsupervised feature selection to extract proximal supervisory signals [21], including seminal work based on spectral graph analysis (e.g., Lap-Score [14], SPEC [45], and NDFS [42]). However, this line of work has mostly been focused on utilizing the topological patterns in the *original feature space* as regularization or constraints for feature selection (e.g., [9, 13, 35]).

It is relatively recent to incorporate *another* network over instances to guide feature selection with the emergence of attributed networks in various domains (e.g., [32, 39]). This provides another source of guidance based on the intuition that the links between instances indicates their similarity in the selected feature space. We contextualize our work in this area. The work most closely related to ours can be generally categorized based on the scope of patterns explored on the network topology into (1) micro-level [22, 40, 41] and (2) macro-level [23, 36]. A common limitation of directly incorporating links and disconnections as *micro-level* guidance is to be susceptible to noisy and incomplete links, which commonly exist in real-world large and complex networks [26]. Our method explores block model as a *macro-level* structural guidance to alleviate such issues.

As for *macro-level (community analysis)*, one of the earliest efforts on network-guided unsupervised feature selection LUFs [36] extracts social dimensions as non-overlapping clusters of nodes to regularize feature selection. NetFS [23] embeds the latent representation learning into the feature selection process, which has been reported to be highly effective on various public datasets. However, the unified optimization framework of NetFS [23] can be influenced by the low-quality features, hence the guidance can deviate from the desirable macro-level network structure. Moreover, their SymNMF-based latent representation learning formulation is consistent with community analysis. Our work is fundamentally different from existing related work both model-wise and methodology-wise.

Block Model is a popular method to group nodes based on structural equivalence/similarity [28], which has been relaxed to stochastic equivalence [1]. It leads to the well-studied area of Stochastic Block Model (SBM), which has long been correlated to community detection [1]. There are studies on the equivalence between NMF-based formulations and SBM at their respective optima [29, 44]. To model disassortative block-level interactions and other challenging structural patterns beyond conventional community structure [2], [11] formulates block model discovery with explicit structural constraints on the image matrix. Recent advances in block model include applications to a wide range of challenging domains, e.g., brain imaging data [4, 5] and Twitter [3]. Methodologically, [27] proposes a framework to efficiently solve block allocation in binary with constraints on the image matrix based on Constrained Programming (CP).

In this paper we use the multiplicative update rules developed by seminal work [7] to build block models from the structural graph. To the best of our knowledge, we are the first to explore modeling the feature selection problem as graph learning regularized by a block model precomputed from the structural graph. There have been many kinds of block models [1] with different characteristics that can potentially match unsupervised feature selection for different linked data. We envision it to facilitate many powerful unsupervised feature selection methods for various complicated real-world attributed networks.

7 CONCLUSION

We propose a novel graph-driven unsupervised feature selection method guided by the block model. The block modeling process in our method is not influenced by the original feature set. Moreover, our similarity graph over nodes in the selected feature space explicitly exploits the relations between nodes. Methodologically, we not only utilize the grouping of nodes as blocks but also require the induced graph to respect the precomputed image matrix. Hence, our method utilizes more complex macro-level network structure beyond conventional community structures. Experiments on various real-world datasets demonstrate the effectiveness of our method as it outperforms baseline methods in finding high-quality features for K-means clustering. For in-depth characterization of our method, we explored the sensitivity of our method regarding hyper-parameters and different block models generated from the graph. This can facilitate the application of our method to domains beyond our exploration. Our method demonstrates the power of using block model to guide unsupervised feature selection. Methodologically, we leave it as future endeavor to jointly learn block model for the graph and select high-quality features.

ACKNOWLEDGMENTS

This research is supported by the National Science Foundation via grant IIS-1910306 and ONR via grant N000141812485.

REFERENCES

- [1] Emmanuel Abbe. 2017. Community detection and stochastic block models: recent developments. *JMLR* 18, 1 (2017), 6446–6531.
- [2] Christopher Aicher, Abigail Z Jacobs, and Aaron Clauset. 2015. Learning latent block structure in weighted networks. *Journal of Complex Networks* 3, 2 (2015).
- [3] Zilong Bai, Buyue Qian, and Ian Davidson. 2018. Discovering models from structural and behavioral brain imaging data. In *SIGKDD*. 1128–1137.
- [4] Zilong Bai, Peter Walker, and Ian Davidson. 2018. Mixtures of block models for brain networks. In *SDM*. 46–54.
- [5] Zilong Bai, Peter Walker, Anna Tschiffely, Fei Wang, and Ian Davidson. 2017. Unsupervised network discovery for brain imaging data. In *SIGKDD*. 55–64.
- [6] Mikhail Belkin and Partha Niyogi. 2002. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NIPS*. 585–591.
- [7] Chris Ding, Tao Li, Wei Peng, and Haesun Park. 2006. Orthogonal nonnegative matrix t-factorizations for clustering. In *SIGKDD*. 126–135.
- [8] Liang Du and Yi-Dong Shen. 2015. Unsupervised feature selection with adaptive structure learning. In *SIGKDD*. 209–218.
- [9] Liang Du, Zhiyong Shen, Xuan Li, Peng Zhou, and Yi-Dong Shen. 2013. Local and global discriminative learning for unsupervised feature selection. In *IEEE ICDM*. 131–140.
- [10] Jennifer G Dy and Carla E Brodley. 2004. Feature selection for unsupervised learning. *JMLR* 5, Aug (2004), 845–889.
- [11] Mohadeseh Ganji, Jeffrey Chan, Peter J Stuckey, James Bailey, Christopher Leckie, Kotagiri Ramamohanarao, and Ian Davidson. 2018. Image constrained blockmodelling: a constraint programming approach. In *SDM*. 19–27.
- [12] Michelle Girvan and Mark EJ Newman. 2002. Community structure in social and biological networks. *PNAS* 99, 12 (2002), 7821–7826.
- [13] Quanquan Gu, Marina Danilevsky, Zhenhui Li, and Jiawei Han. 2012. Locality preserving feature learning. In *AISTATS*. 477–485.
- [14] Xiaofei He, Deng Cai, and Partha Niyogi. 2006. Laplacian score for feature selection. In *NIPS*. 507–514.
- [15] Chenping Hou, Feiping Nie, Dongyun Yi, and Yi Wu. 2011. Feature selection via joint embedding learning and sparse regression. In *IJCAI*.
- [16] Xiao Huang, Jundong Li, and Xia Hu. 2017. Label Informed Attributed Network Embedding. In *WSDM*. 731–739.
- [17] Xiao Huang, Qingquan Song, Jundong Li, and Xia Hu. 2018. Exploring expert cognition for attributed network embedding. In *WSDM*. 270–278.
- [18] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *ICLR* (2016).
- [19] Ron Kohavi, George H John, et al. 1997. Wrappers for feature subset selection. (1997).
- [20] Solomon Kullback. 1997. *Information theory and statistics*. Courier Corporation.
- [21] Jundong Li, Kewei Cheng, Suhang Wang, Fred Morstatter, Robert P Trevino, Jiliang Tang, and Huan Liu. 2017. Feature selection: A data perspective. *ACM Computing Surveys (CSUR)* 50, 6 (2017), 1–45.
- [22] Jundong Li, Ruocheng Guo, Chenghao Liu, and Huan Liu. 2019. Adaptive unsupervised feature selection on attributed networks. In *SIGKDD*. 92–100.
- [23] Jundong Li, Xia Hu, Liang Wu, and Huan Liu. 2016. Robust unsupervised feature selection on networked data. In *SDM*. 387–395.
- [24] Zechao Li, Yi Yang, Jing Liu, Xiaofang Zhou, and Hanqing Lu. 2012. Unsupervised feature selection using nonnegative spectral analysis. In *AAAI*.
- [25] Chih-Jen Lin. 2007. Projected gradient methods for nonnegative matrix factorization. *Neural computation* 19, 10 (2007), 2756–2779.
- [26] Huan Liu, Fred Morstatter, Jiliang Tang, and Reza Zafarani. 2016. The good, the bad, and the ugly: uncovering novel research opportunities in social media mining. *International Journal of Data Science and Analytics* 1, 3–4 (2016), 137–143.
- [27] Alex Matteniet, Ian Davidson, Siegfried Nijssen, and Pierre Schaus. 2019. Generic Constraint-Based Block Modeling Using Constraint Programming. In *CP*. Springer, 656–673.
- [28] Berndt Müller, Joachim Reinhardt, and Michael T Strickland. 2012. *Neural networks: an introduction*. Springer Science & Business Media.
- [29] Subhadeep Paul and Yuguo Chen. 2016. Orthogonal symmetric non-negative matrix factorization under the stochastic block model. *arXiv preprint arXiv:1605.05349* (2016).
- [30] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *JMLR* 12 (2011), 2825–2830.
- [31] Mingjie Qian and Chengxiang Zhai. 2013. Robust unsupervised feature selection. In *IJCAI*.
- [32] Nahid Safari-Alighiarloo, Mohammad Taghizadeh, Mostafa Rezaei-Tavirani, Bahram Goliaei, and Ali Asghar Peyvandi. 2014. Protein-protein interaction networks (PPI) and complex diseases. *Gastroenterol Hepatol Bed Bench* (2014).
- [33] Patricia Iglesias Sánchez, Emmanuel Müller, Uwe Leo Korn, Klemens Böhm, Andrea Kappes, Tanja Hartmann, and Dorothea Wagner. 2015. Efficient algorithms for a robust modularity-driven clustering of attributed graphs. In *SDM*. 100–108.
- [34] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
- [35] Lei Shi, Liang Du, and Yi-Dong Shen. 2014. Robust spectral learning for unsupervised feature selection. In *IEEE ICDM*. 977–982.
- [36] Jiliang Tang and Huan Liu. 2012. Unsupervised feature selection for linked social media data. In *SIGKDD*. 904–912.
- [37] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale information network embedding. In *WWW*. 1067–1077.
- [38] Lei Tang and Huan Liu. 2009. Relational learning via latent social dimensions. In *SIGKDD*. 817–826.
- [39] Io Taxisidou and Peter M Fischer. 2014. Online analysis of information diffusion in twitter. In *WWW*. 1313–1318.
- [40] Xiaokai Wei, Bokai Cao, and S Yu Philip. 2016. Unsupervised feature selection on networks: a generative view. In *AAAI*.
- [41] Xiaokai Wei, Sihong Xie, and Philip S Yu. 2015. Efficient partial order preserving unsupervised feature selection on networks. In *SDM*. 82–90.
- [42] Yi Yang, Heng Tao Shen, Zhigang Ma, Zi Huang, and Xiaofang Zhou. 2011. L2, 1-norm regularized discriminative feature selection for unsupervised. In *IJCAI*.
- [43] Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. 2018. Network representation learning: A survey. *IEEE transactions on Big Data* (2018).
- [44] Zhong-Yuan Zhang, Yujie Gai, Yu-Fei Wang, Hui-Min Cheng, and Xin Liu. 2018. On equivalence of likelihood maximization of stochastic block model and constrained nonnegative matrix factorization. *Physica A: Statistical Mechanics and its Applications* 503 (2018), 687–697.
- [45] Zheng Zhao and Huan Liu. 2007. Spectral feature selection for supervised and unsupervised learning. In *ICML*. 1151–1157.
- [46] Chen Zhe, Aixin Sun, and Xiaokui Xiao. 2019. Community detection on large complex attribute network. In *SIGKDD*. 2041–2049.

A PROOF FOR THEOREM 3.1

THEOREM (LEAST SQUARES OPTIMAL $\mathbf{M}$ IN CLOSED FORM). *Given $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{F} \in \{0, 1\}^{n \times k}$. If $\mathbf{D} = \mathbf{F}^T \mathbf{F}$ is a diagonal matrix with positive diagonal elements, then*

$$\underset{\mathbf{X}}{\operatorname{argmin}} \|\mathbf{A} - \mathbf{F}\mathbf{X}\mathbf{F}^T\|_F^2 = \mathbf{D}^{-1} \mathbf{F}^T \mathbf{A} \mathbf{F} \mathbf{D}^{-1} \quad (16)$$

PROOF. Let $\mathbf{M} = \underset{\mathbf{X}}{\operatorname{Argmin}} \|\mathbf{A} - \mathbf{F}\mathbf{X}\mathbf{F}^T\|_F^2$. Then the following equations exist by vectorizing matrices $\mathbf{M}$ and $\mathbf{A}$:

$$\begin{aligned} \operatorname{vec}(\mathbf{M}) &= \underset{\mathbf{x}}{\operatorname{argmin}} \|\operatorname{vec}(\mathbf{A}) - (\mathbf{F} \otimes \mathbf{F})\mathbf{x}\|_2^2 \\ &= [(\mathbf{F} \otimes \mathbf{F})^T (\mathbf{F} \otimes \mathbf{F})]^{-1} (\mathbf{F} \otimes \mathbf{F})^T \operatorname{vec}(\mathbf{A}) \\ &= [(\mathbf{F}^T \mathbf{F}) \otimes (\mathbf{F}^T \mathbf{F})]^{-1} (\mathbf{F} \otimes \mathbf{F})^T \operatorname{vec}(\mathbf{A}) \\ &= [\mathbf{D} \otimes \mathbf{D}]^{-1} (\mathbf{F} \otimes \mathbf{F})^T \operatorname{vec}(\mathbf{A}) \\ &\stackrel{(*)}{=} [\mathbf{D}^{-1} \otimes \mathbf{D}^{-1}] (\mathbf{F}^T \otimes \mathbf{F}^T) \operatorname{vec}(\mathbf{A}) \\ &= [(\mathbf{D}^{-1} \mathbf{F}^T) \otimes (\mathbf{D}^{-1} \mathbf{F}^T)] \operatorname{vec}(\mathbf{A}) \end{aligned} \quad (17)$$

where equation (*) sustains because $\mathbf{D} = \mathbf{F}^T \mathbf{F}$ is a diagonal matrix with positive diagonal elements, hence $\mathbf{D}$ is non-singular. Lastly, rewrite $\mathbf{M}$ and $\mathbf{A}$ into matrices:

$$\mathbf{M} = \mathbf{D}^{-1} \mathbf{F}^T \mathbf{A} \mathbf{F} \mathbf{D}^{-1} \quad (18)$$

□

B DETAILS ON EXPERIMENTS

B.1 Datasets Preprocessing

- **BlogCatalog** [16, 38] BlogCatalog is a social network dataset, which contains an undirected graph of 5196 nodes and 171743 edges, where each node is a BlogCatalog user, and each edge indicates whether the two users are friends on the website. The attributes are the appearance of 8189 keywords in the blog descriptions. We obtain the dataset from <http://people.tamu.edu/~xhuang/BlogCatalog.mat.zip>.
- **Citeseer** [34] Citeseer is a citation network dataset, which contains a directed graph of 3327 nodes where each edge indicates a citation of a paper to another, as well as the appearance of 3703 words in 3312 papers out of 3327 papers in the graph. We obtain the dataset from <https://github.com/tkipf/gcn> [18], and we cross-check the data set with the data set obtained from <https://linqs-data.soe.ucsc.edu/public/lbc/citeseer.tgz>. The citation network in the dataset is a directed graph. We transform the graph to an undirected graph by creating (u, v) and (v, u) edges for every edge (u, v) in the directed graph. There are 15 papers that do not have attributes, we remove the corresponding nodes and edges from the graph. As a result, the Citeseer dataset we use in the experiment contains a citation graph of 3312 nodes and 4660 undirected edges.
- **Cora** [34] Cora is a citation network dataset, which contains a directed graph of 2708 nodes where each edge indicates a citation of a paper to another, and appearance of 1433 words of all papers in the graph. We obtain the dataset from <https://github.com/tkipf/gcn> [18], and we cross-check the data set with the one obtained from <https://linqs-data.soe>.

[ucsc.edu/public/lbc/cora.tgz](https://linqs-data.soe.ucsc.edu/public/lbc/cora.tgz). The citation network in the dataset is a directed graph. We transform the graph to an undirected graph by creating (u, v) and (v, u) edges for every edge (u, v) in the directed graph. As a result, the Cora dataset we use in the experiment contains a citation graph of 2708 nodes and 5278 undirected edges.

B.2 K-means Clustering

For our experiments, we use scikit-learn's implementation of K-means clustering algorithm. We also use sklearn's normalize function to normalize the feature sets before running the clustering algorithm [30]. For all experiments involving K-means clustering, we run the algorithm 20 times to compensate for random initialization.

B.3 Definitions of Metrics for Clustering Performance Evaluation

Accuracy (ACC)

$$\text{ACC} = \frac{\sum_{i \in [n]} \mathbb{I}(l_i = \sigma(c_i))}{n} \quad (19)$$

where c_i is the clustering result of data point i and l_i is its ground-truth class label. Permutation function σ maps c_i to a class label using Kuhn-Munkres Algorithm. We utilize the implementation of ACC from https://github.com/Tony607/Keras_Deep_Clustering/blob/master/metrics.py in our experiments.

Normalized Mutual Information (NMI)

$$\text{NMI}(\mathcal{L}, C) = \frac{\text{MI}(\mathcal{L}, C)}{\max(H(\mathcal{L}), H(C))} \quad (20)$$

where $H(\mathcal{L})$ and $H(C)$ are respectively the entropy of $\mathcal{L}$ (the node/instance grouping based on class labels) and C (the instance clustering based on selected features) respectively.

B.4 Baseline Methods: Links to Source Codes and Setting Hyper-parameters

- **MMPOP** [41]
 - **Link to Source Code:** http://www.cse.lehigh.edu/~sxie/codes/optimization_pop.py.
 - **Setting parameters:** We keep the default parameters across all experiments.
- **NetFS** [23]
 - **Link to Source Code:** <http://people.virginia.edu/~jl6qk/code/NetFS.zip>.
 - **Setting parameters:** We set α to be 10 and β to be 0.1 according to the original paper's experiments.
- **NDFS** [24]
 - **Link to Source Code:** We use skfeature [21]'s implementation of Nonnegative Discriminative Feature Selection algorithm from https://github.com/jundongli/scikit-feature/blob/master/skfeature/example/test_NDFS.py
 - **Setting parameters:** We utilize the default setting in skfeature's provided example. The number of clusters is set according to the number of distinct ground truth labels in each dataset. The number of neighbors is set to 5 as suggested in the original paper.

- **SPEC** [45]
 - **Link to Source Code:** We use skfeature [21]’s implementation of Spectral Feature Selection algorithm from https://github.com/jundongl/scikit-feature/blob/master/skfeature/example/test_SPEC.py
 - **Setting parameters:** We utilize the default setting in skfeature’s provided example. The number of clusters is set according to the number of distinct ground truth labels in each dataset.
- **Laplacian Score** [14]
 - **Link to Source Code:** We use skfeature [21]’s implementation of Laplacian Score feature selection algorithm from https://github.com/jundongl/scikit-feature/blob/master/skfeature/example/test_lap_score.py
 - **Setting parameters:** We utilize the default setting in the skfeature’s provided example. The number of clusters is set according to the number of distinct ground truth labels in each dataset. The number of neighbors is set to 5 as suggested in the original paper.