
On The Projection Operator to A Three-view Cardinality Constrained Set

Haichuan Yang¹ Shupeng Gui¹ Chuyang Ke¹ Daniel Stefankovic¹ Ryohei Fujimaki² Ji Liu¹

Abstract

The cardinality constraint is an intrinsic way to restrict the solution structure in many domains, for example, sparse learning, feature selection, and compressed sensing. To solve a cardinality constrained problem, the key challenge is to solve the projection onto the cardinality constraint set, which is NP-hard in general when there exist multiple overlapped cardinality constraints. In this paper, we consider the scenario where the overlapped cardinality constraints satisfy a Three-view Cardinality Structure (TVCS), which reflects the natural restriction in many applications, such as identification of gene regulatory networks and task-worker assignment problem. We cast the projection into a linear programming, and show that for TVCS, the vertex solution of this linear programming is the solution for the original projection problem. We further prove that such solution can be found with the complexity proportional to the number of variables and constraints. We finally use synthetic experiments and two interesting applications in bioinformatics and crowdsourcing to validate the proposed TVCS model and method.

1. Introduction

The cardinality constraint is an intrinsic way to restrict the solution structure in many real problems, for example, sparse learning (Olshausen & Field, 1997), feature selection (Zhang, 2009), and compressed sensing (Candes et al., 2006). The generic cardinality constrained optimization can be expressed as

$$\min_{\mathbf{w} \in \mathbb{R}^p} f(\mathbf{w}) \quad (1a)$$

$$\text{subject to } \|\mathbf{w}_g\|_0 \leq \mathbf{s}^g \quad \forall g \in \mathcal{G} \quad (1b)$$

¹University of Rochester, Rochester, NY, USA ²NEC, Cupertino, CA, USA. Correspondence to: Haichuan Yang <h.yang@rochester.edu>.

where \mathbf{w} is the optimization variable, g is an index subset of $[p] := \{1, 2, \dots, p\}$, \mathbf{w}_g is the sub-vector of \mathbf{w} indexed by g . $\|\mathbf{w}_g\|_0$ denotes the cardinality of the sub-vector, i.e., the number of nonzeros in \mathbf{w}_g , \mathcal{G} is the hyper set of all predefined groups, and $\mathbf{s} \in \mathbb{R}^{|\mathcal{G}|}$ is the upper bound vector - $\mathbf{s}^g \in \mathbb{R}$ refers to the upper bound of the sparsity over group g . Objective f is the loss function which could be defined with different form according to the specific application. The problem (1) refers to a nonconvex optimization (NP-hard) due to the cardinality constraint. Some efficient iterative methods such as IHT (Yuan et al., 2014), CoSaMP (Needell & Tropp, 2009), GradMP (Nguyen et al., 2012), and their variants can guarantee to solve the original problem under some mild conditions. A key component in all of these methods is the projection operator

$$\mathbf{P}_{\Omega(\mathcal{G}, \mathbf{s})}(\mathbf{v}) := \arg \min_{\mathbf{w} \in \Omega(\mathcal{G}, \mathbf{s})} \|\mathbf{w} - \mathbf{v}\|^2 \quad (2)$$

where $\Omega(\mathcal{G}, \mathbf{s})$ denotes the feasible set to the constraint (1b). While in some special case, for example, $\mathcal{G} = \{[p]\}$, the projection is trivial, it is quite challenging, especially when \mathcal{G} includes multiple *overlapped* index sets (even NP-hard in some cases).

In this paper, we consider the scenario where the overlapped cardinality constraints (1b) satisfy a Three-view Cardinality Structure (TVCS):

Definition 1. (Three-view Cardinality Structure (TVCS)) For $\Omega(\mathcal{G}, \mathbf{s})$, the hyper set \mathcal{G} consisting of subsets of $[p]$ admits the TVCS structure if the following conditions are satisfied:

- There exists a partition $\mathcal{G}_0, \mathcal{G}_1$ and \mathcal{G}_2 such that $\mathcal{G} = \mathcal{G}_0 \cup \mathcal{G}_1 \cup \mathcal{G}_2$;
- $\mathcal{G}_0 = \{[p]\}$;
- All element sets in \mathcal{G}_1 have no overlap;
- All element sets in \mathcal{G}_2 have no overlap.

This definition basically requires that \mathcal{G} can be partitioned into three hyper sets $\mathcal{G}_0, \mathcal{G}_1$, and \mathcal{G}_2 , and overlaps can only happen between element sets in different hyper sets. \mathcal{G}_0 is

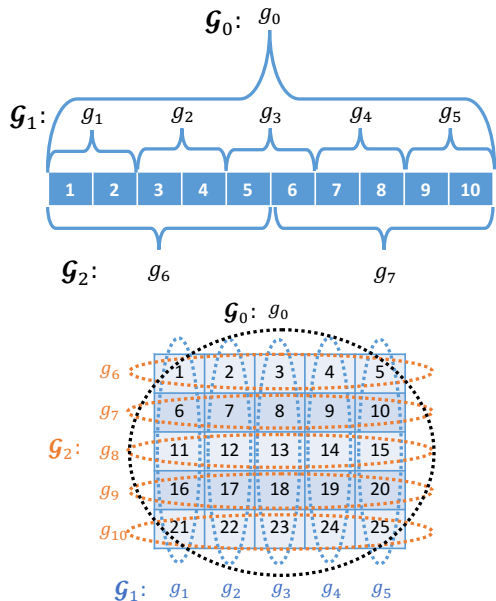


Figure 1. Two examples of groups for TVCS model. The first model has $p = 10$, $\mathcal{G}_0 = \{g_0\}$, $\mathcal{G}_1 = \{g_1, g_2, g_3, g_4, g_5\}$, $\mathcal{G}_2 = \{g_6, g_7\}$. The second model organizes elements as matrix ($p = 25$). Each row and column is a group, and $\mathcal{G}_0 = \{g_0\}$, $\mathcal{G}_1 = \{g_1, g_2, g_3, g_4, g_5\}$, $\mathcal{G}_2 = \{g_6, g_7, g_8, g_9, g_{10}\}$.

usually used to restrict the overall sparsity. Figure 1 provides two examples of \mathcal{G} for TVCS.

The TVCS model is motivated from some important applications, for example, in recommendation, task-worker assignment, and bioinformatics.

- **Online recommendation.** Suppose we want to recommend a certain number of books (among p books) to a customer - corresponding to the \mathcal{G}_0 based sparsity constraint. Among the selected books, we want to maintain some diversities - the recommended books by the same author should not exceed a certain number (\mathcal{G}_1 based sparsity constraint) and about the same topic should not exceed a certain number either (\mathcal{G}_2 based sparsity constraint). One can refer to the top graph in Figure 1: \mathcal{G}_1 is grouped by authors and \mathcal{G}_2 is grouped by topics.
- **Task-worker assignment.** Suppose we have a bunch of tasks and workers, and we want to assign the tasks to workers. For example, in crowdsourcing, we usually assign several different workers to each task since we want to use the answers from multiple workers to improve the accuracy. On the other hand, each worker is usually assigned to multiple tasks so there is a “many to many” relationship in this assignment.

The goal is to pursue the optimal assignment under a certain criteria in crowdsourcing, while satisfying some restrictions. For example, the total assignments should be bounded by the total budget (corresponding to \mathcal{G}_0), the total cost of assignments to a single worker cannot exceed a certain threshold (corresponding to \mathcal{G}_1), and the total cost of assignments on a single task cannot exceed a certain threshold (corresponding to \mathcal{G}_2). Let X be the assignment matrix, and its rows are indexed by workers and the columns are indexed by tasks. These constraints can be illustrated by the bottom graph in Figure 1.

- **Identification of gene regulatory networks.** The essential goal of identifying gene regulatory network is to identify a weighted directed graph, which can be represented by a square matrix W with $p = N \times N$ elements in total where N is the number of vertices. A sparse network constraint is to restrict the in-degree and out-degree for each vertex, which corresponds to the sparsity in each row and column of W .

To solve the TVCS constrained projection (2), we show an interesting connection between the projection and a linear programming (LP) that the vertex solution to this linear programming is an integer solution which solves the original problem.

To find an integer solution to such LP efficiently, we formulate it into a feasibility problem, and further an equivalent quadratic convex optimization. By using the rounding technique, we can avoid solving the exact solution of this LP problem. We propose an iterative algorithm to solve it and each iteration can be completed in linear time. We also show that the iterate linearly converges to the optimal point. Finally, the proposed TVCS model is validated by the synthetic experiment and two important and novel applications in identification of gene regulatory networks and task assignment problem of crowdsourcing.

2. Related Works

Recent years have witnessed many research works in the field of structured sparsity and group-based sparsity. Yuan & Lin (2006) introduced the group LASSO, which pursues group-wise sparsity that restricts the number of groups for the selected variables. Jenatton et al. (2011) construct a hierarchical structure over the variables and use group LASSO with overlapped groups to solve it. Exclusive LASSO (Zhou et al., 2010; Kong et al., 2014) was proposed for the exclusive group sparsity which can be treated as relaxing our cardinality constraints to convex regularizations. In (Kong et al., 2014), the authors discussed the overlapping situation and tried to solve the problem using convex relaxation, which is different from our approach.

Besides the aforementioned works, some proposed more general models to cover various sparsity structures. Bach et al. (2012) extended the usage of L_1 -norm relaxation to several different categories of structures. And recently, another generalization work (El Halabi & Cevher, 2015) proposed convex envelopes for various sparsity structures. They built the framework by defining a totally unimodular penalty, and showed how to formulate different sparsity structures using the penalty. The work above concentrated on using convex relaxation to control the sparsity.

Besides using convex relaxation, there are also some works focusing on projection-based methods. When the exact projection operator was provided, Baraniuk et al. (2010) extended the traditional IHT and CoSaMP methods to general sparsity structures. In this work, the authors also introduced the projection operator for block sparsity and tree sparsity. Cevher et al. (2009) investigated cluster sparsity and they applied dynamic programming to solve the projection operator for their sparsity model. Hegde et al. (2009) introduced a “spike trains” signal model, which is also related to exclusive group sparsity. Its groups always have consecutive coordinates, and each group cannot contain more than one nonzero element. To solve the projection problem of their model, they showed the basic feasible solutions of the relaxed linear programming (LP) are always integer points. In our work, we also use LP to solve the projection problem. But our model defines the group structure differently and aims at different applications.

In addition, there are some works for the cases without an efficient exact projection operator (Hegde et al., 2015a;b; Nguyen et al., 2014). This is meaningful since the projection operator for complex structured sparsity often involves solving complicated combinatorial optimization problems. Hegde et al. (2015a) discussed how to guarantee convergence if using approximate projection in IHT and CoSaMP for compressive sensing. They proved that the convergence needs a “head approximation” to project the update (gradient) before applying it. Hegde et al. (2015b) proposed a general framework to formulate a series of models as a weighted graph, and designed an efficient approximate projection operator for the models. Nguyen et al. (2014) applied the approximate projection-based IHT and CoSaMP to general convex functions and stochastic gradients.

3. Preliminary: GradMP and IHT Frameworks

This section briefly reviews two commonly used algorithm frameworks to solve the cardinality constrained optimization (1): iterative hard thresholding (IHT) (Yuan et al., 2014; Nguyen et al., 2014) and gradient matching pursuit (GradMP) (Nguyen et al., 2012; 2014) (the general version of CoSaMP (Needell & Tropp, 2009)) for solving

cardinality constrained problem. Other methods like hard thresholding pursuit (HTP) also follows similar steps and has been shown to be effective both empirically and theoretically (Yuan et al., 2016). The procedures of IHT and GradMP for our model are shown in Algorithms 1 and 2, where $\text{supp}(\cdot)$ is the support set of the argument vector.

Therefore, one can see that the efficiency of both algorithms relies on the computation of the gradient and the projection. To avoid the expensive computation of the gradient, GradMP and IHT can be extended to the stochastic versions (Nguyen et al., 2014) by assigning \mathbf{g} the stochastic gradient at the gradient computation step.

Both Algorithms 1 and 2 (and their stochastic variants) guarantee some nice properties: the iterate converges to a small ball surrounding the true solution at a linear rate under certain RIP-type conditions (Nguyen et al., 2014) and the radius of such ball converges to zero when the number of samples goes to infinity.

Algorithm 1: Iterative Hard Thresholding.

Input: Sparsity parameter s .

Result: Problem solution \mathbf{w}^t .

Initialize $\mathbf{w}^0, t = 0$;

while *stop criterion is not met* **do**

$\mathbf{g} = \nabla f(\mathbf{w}^t)$; // Gradient computation

$\mathbf{z}^t = \mathbf{w}^t - \gamma \mathbf{g}$; // Gradient descent

$\mathbf{w}^{t+1} = \mathbf{P}_{\Omega(\mathcal{G}, s)}(\mathbf{z}^t)$; // Projection

$t = t + 1$;

end

Algorithm 2: Gradient Matching Pursuit.

Input: Sparsity parameter s .

Result: Problem solution \mathbf{w}^t .

Initialize $\mathbf{w}^0, t = 0$;

while *stop criterion is not met* **do**

$\mathbf{g} = \nabla f(\mathbf{w}^t)$; // Gradient computation

$\Gamma = \text{supp}(\mathbf{P}_{\Omega(\mathcal{G}, 2s)}(\mathbf{g}))$;

$\hat{\Gamma} = \Gamma \cup \text{supp}(\mathbf{w}^t)$; // Subspace selection

$\mathbf{z}^t = \arg \min_{\text{supp}(\mathbf{z})=\hat{\Gamma}} f(\mathbf{z})$; // Subspace optimization

$\mathbf{w}^{t+1} = \mathbf{P}_{\Omega(\mathcal{G}, s)}(\mathbf{z}^t)$; // Projection

$t = t + 1$;

end

A common component in Algorithms 1 and 2 is the projection operator. If all the groups except $[p]$ in \mathcal{G} do not overlap each other, the projection problem can be easily solved by sequential projections (Yang et al., 2016). But for those cases involving overlapped groups, it is generally challenging to solve them efficiently.

4. Projection Operator

This section introduces how to solve the essential projection step. Note that the projection onto a nonconvex set is NP-hard in general. By utilizing the special structure of TVCS, we show that the projection can be solved efficiently. Due to the page limitation, all proofs are provided in the supplementary material.

4.1. LP Relaxation

Firstly, we can cast the projection problem (2) to an equivalent integer linear programming problem (ILP) according to Lemma 1.

Lemma 1. *The projection problem (2) is equivalent to the following integer linear programming problem (ILP):*

$$\begin{aligned} \max_{\mathbf{x}} \quad & \langle \mathbf{v}^2, \mathbf{x} \rangle \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{s} \\ & \mathbf{x} \in \{0, 1\}^p \end{aligned} \quad (3)$$

where \mathbf{v}^2 is applying element-wise square operation on vector \mathbf{v} . A is a $|\mathcal{G}| \times p$ matrix which is defined as:

$$A = \begin{bmatrix} \mathbf{1}^\top \\ C \end{bmatrix} \quad (4)$$

where $C \in \{0, 1\}^{|\mathcal{G}_1 \cup \mathcal{G}_2| \times p}$, whose rows represent the indicator vector of each group $g \in \mathcal{G}_1$ and \mathcal{G}_2 .

Each row in A corresponds to one group g from \mathcal{G} . For example, $C_{ij} = 1$ if the j -th coordinate is in the i -th group, otherwise $C_{ij} = 0$. The first row $\mathbf{1}^\top$ corresponds to the overall sparsity i.e. \mathcal{G}_0 .

It is NP-hard to solve an ILP in general. One common way to handle such ILP is making a linear programming (LP) relaxation. In our case, we can use a box constraint $\mathbf{x} \in [0, 1]^p$ to replace the integer constraint $\mathbf{x} \in \{0, 1\}^p$:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \langle \mathbf{v}^2, \mathbf{x} \rangle \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{s} \\ & \mathbf{x} \in [0, 1]^p \end{aligned} \quad (5)$$

However, there is no guarantee that a general ILP can be solved via its LP relaxation, because the solution of the relaxed LP is not always integer. Although one can make a rounding to the LP solution and acquire a integer solution, such solution is not guaranteed to be optimal (or even feasible) to the original ILP.

Fortunately, due to the special structure of our TVCS model, we find that its relaxed LP has some nice properties which make it possible to get the optimal solution of the ILP efficiently. The following theorem reveals the relationship between the ILP problem and the relaxed LP problem.

Theorem 2. *Given \mathcal{G} satisfying TVCS, all the vertices of the feasible set to (5) are integer points. Furthermore, there is an optimal solution on the vertex that solves the ILP (3).*

This theorem suggests that finding a vertex solution of the relaxed LP can solve the original projection problem onto a TVCS \mathcal{G} . The proof basically shows that matrix A (for TVCS) is a totally unimodular matrix (Papadimitriou & Steiglitz, 1982). We provide the detailed proof in the supplementary material.

4.2. Linearly Convergent Algorithm for Projection Operator onto TVCS

To find a solution on the vertex, one can use the Simplex method. Although Simplex method guarantees to find an optimal solution on the vertex and could be very efficient in practice, it does not have a deterministic complexity bound. In the IHT and GradMP algorithms, projection operator is only a sub-procedure in one iteration. Hence, we are usually supposed to solve lots of instances of problem (3). Simplex might be efficient practically, but its worst case may lead to exponential time complexity (Papadimitriou & Steiglitz, 1982). In this section, the integer solution to the linear programming can be found within the complexity proportional to the number of variables and constraints.

Equivalent Feasibility Problem Formulation. The dual of LP problem (5) can be written as:

$$\begin{aligned} \min_{\mathbf{y}} \quad & \langle [\mathbf{s}^\top \ \mathbf{1}^\top]^\top, \mathbf{y} \rangle \\ \text{subject to} \quad & [A^\top \ I] \mathbf{y} \geq \mathbf{v}^2, \mathbf{y} \geq \mathbf{0} \end{aligned} \quad (6)$$

Since the duality gap of LP is zero, combining the primal LP (5) and dual LP (6), we can formulate an equivalent problem, i.e. the feasibility problem over the following constraints:

$$\begin{aligned} \text{find} \quad & \mathbf{x}, \mathbf{y} \\ \text{subject to} \quad & \langle [\mathbf{s}^\top \ \mathbf{1}^\top]^\top, \mathbf{y} \rangle = \langle \mathbf{v}^2, \mathbf{x} \rangle \\ & [A^\top \ I] \mathbf{y} \geq \mathbf{v}^2 \\ & \begin{bmatrix} A \\ I \end{bmatrix} \mathbf{x} \leq \begin{bmatrix} \mathbf{s} \\ \mathbf{1} \end{bmatrix} \\ & \mathbf{y} \geq \mathbf{0}, \mathbf{x} \geq \mathbf{0} \end{aligned}$$

Iterative Algorithm. The feasibility problem with linear constraints above is equivalent to the following optimiza-

tion problem:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{y}} \quad & \frac{1}{2} (\langle [\mathbf{s}^\top \mathbf{1}^\top]^\top, \mathbf{y} \rangle - \langle \mathbf{v}^2, \mathbf{x} \rangle)^2 \\ & + \frac{1}{2} \|\mathbf{v}^2 - [A^\top I] \mathbf{y}\|_+^2 + \frac{1}{2} \|[A\mathbf{x} - \mathbf{s}]\|_+^2 \end{aligned} \quad (7)$$

subject to $\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}, \mathbf{y} \geq \mathbf{0}$

where $[\mathbf{z}]_+$ is the element-wise hinge operator, i.e. it transforms each element \mathbf{z}_i to $\max(\mathbf{z}_i, 0)$.

This is a convex optimization problem with a quadratic objective and box constraints. We adopt the projected gradient descent to solve this problem, and show it converges linearly.

Theorem 3. *For the optimization problem with the form*

$$\begin{aligned} \min_{\mathbf{z}} \quad & f(\mathbf{z}) := \|[A\mathbf{z} - \mathbf{a}]\|_+^2 + \|B\mathbf{z} - \mathbf{b}\|^2 \\ \text{subject to} \quad & \mathbf{z} \in \Omega \end{aligned}$$

where $\Omega = \{\mathbf{z} \mid C\mathbf{z} \leq \mathbf{c}\}$, the projected gradient descent algorithm $\mathbf{z}^{t+1} \leftarrow P_\Omega(\mathbf{z}^t - \gamma \nabla f(\mathbf{z}^t))$ has a linear convergence rate with some $\alpha < 1$ (depending on A and B):

$$\|\mathbf{z}^{t+1} - P_{\mathbf{z}^*}(\mathbf{z}^{t+1})\| \leq \alpha \|\mathbf{z}^t - P_{\mathbf{z}^*}(\mathbf{z}^t)\|,$$

where $P_{\mathbf{z}^*}(\cdot)$ is the projection onto the optimal solution set.

Notice that the objective function f in Theorem 3 is not necessarily strongly convex, which means the well recognized linear convergence conclusion from the strong convexity is not applicable here.

Theorem 3 mainly applies Hoffman's Theorem (Hoffman, 2003) to show that f is an optimal strongly convex function (Liu & Wright, 2015). This leads to a linear convergence rate.

The convergence rate $\alpha = 1/(1+\frac{\lambda}{L})$, where λ is the Hoffman constant (Hoffman, 2003) that depends on A, B and is always positive. L is the Lipschitz continuous gradient constant. More details are included in the supplementary materials.

To show the complexity of this algorithm, we firstly count how many iterations we need. Since we know that we can just make a rounding¹ to the result \mathbf{x}^t when we attain $\|\mathbf{x}^t - \tilde{\mathbf{x}}^*\|_\infty < 0.5$. Let $\mathbf{z} := [\mathbf{x}^\top \mathbf{y}^\top]^\top$ represent all the variables in (7). Because $\|\mathbf{z}^t - \mathbf{z}^*\| \geq \|\mathbf{z}^t - \mathbf{z}^*\|_\infty \geq \|\mathbf{x}^t - \mathbf{x}^*\|_\infty$, we can do the rounding safely when $\|\mathbf{z}^t - \mathbf{z}^*\| < 0.5$, where $\mathbf{z}^*, \mathbf{x}^*$ are the optimal points of this problem. According to

¹Acute readers may notice that the convergent point may be on the face of the polytope in some cases instead of vertex. However, we can add a small random perturbation to ensure the optimal point to be vertices with probability 1.

Theorem 3, we have the linear convergence rate $\alpha < 1$, so the number of iterations we need is

$$t > \log_{\alpha} \frac{1}{2\|\mathbf{z}^0 - \mathbf{z}^*\|}$$

Therefore, we claim that we can obtain the solution \mathbf{x}^* by rounding after $\log_{\alpha} \frac{1}{2\|\mathbf{z}^0 - \mathbf{z}^*\|}$ iterations.

Secondly, we show that the computation complexity in each iteration is linear with dimensionality p and the amount of groups $|\mathcal{G}|$. Since each column of A contains at most 3 nonzero elements, the complexity of the matrix multiplications in computing the gradient of (7) is $O(p + |\mathcal{G}|)$. Together with other computation, the complexity for each iteration is $O(p + |\mathcal{G}|)$.

5. Empirical Study

This section will validate the proposed method on both synthetic data and two practical applications: crowdsourcing and identification of gene regulatory networks.

5.1. Linear Regression and Classification on Synthetic Data

In this section, we validate the proposed method with linear regression objective and squared hinge objective (classification) on synthetic data. Let $\mathbf{w} \in \mathbb{R}^{\sqrt{p} \times \sqrt{p}}$ be a matrix, \mathcal{G}_1 and \mathcal{G}_2 are defined as groups with all rows and all columns respectively. The linear regression loss is defined as $\sum_{i=1}^n (\langle X_i, \mathbf{w} \rangle - y_i)^2$ and the squared hinge loss is defined as $\sum_{i=1}^n \max(0, 1 - y_i \langle X_i, \mathbf{w} \rangle)^2$, where n is the total number of training samples. X_i and y_i are the features and label of the i -th sample respectively.

In the linear regression experiment, the true model $\bar{\mathbf{w}} \in \mathbb{R}^{\sqrt{p} \times \sqrt{p}}$ is generated from the following procedure: generate a random vector and apply the projection operator to get a support set which satisfy our sparsity constraints; the elements of positions in support set are drawn from standard normal distribution. p is fixed as 400 and n is gradually increased. The group sparsity upper bounds \mathbf{s}^g for $g \in \mathcal{G}_1$ and $g \in \mathcal{G}_2$ are uniformly generated from the integers in the range $[1, \sqrt{p}]$. The overall sparsity upper bound is set by $0.8 \times \min(\sum_{g \in \mathcal{G}_1} \mathbf{s}^g, \sum_{g \in \mathcal{G}_2} \mathbf{s}^g)$. Each X_i 's is an $\sqrt{p} \times \sqrt{p}$ i.i.d. Gaussian random matrix. y_i is generated from $y_i = \langle X_i, \bar{\mathbf{w}} \rangle + e_i$, where e_i is the i.i.d. Gaussian random noise drawn from $\mathcal{N}(0, 0.01^2)$. We compare the proposed methods to bilevel exclusive sparsity with non-overlapped groups (row-wise or column-wise) (Yang et al., 2016), overall sparsity (Needell & Tropp, 2009), and exclusive LASSO (Kong et al., 2014). For fairness we project the final result of all the compared methods to satisfy all constraints. All the experiments are repeated 30 times and we use the averaged result. We use selection recall and successful recovery rate to evaluate the performance. Selection

recall is defined as $|\text{supp}(\mathbf{w}^*) \cap \text{supp}(\bar{\mathbf{w}})| / \|\bar{\mathbf{w}}\|_0$, where \mathbf{w}^* is the optimization result. Successful recovery rate is the ratio of the successful feature selection i.e. $\text{supp}(\mathbf{w}^*) = \text{supp}(\bar{\mathbf{w}})$ to the total number of repeated experiments. In Figure 2 we can observe that our model with all sparsity constraints always have the best performance. While the performance of exclusive LASSO and our method is comparable when the number of samples are very limited, our method outperforms exclusive LASSO when the number of samples increases.

For classification experiments, we use the same settings of sparsity with linear regression. Here we set $p = 400$, and change n from 200 to 800. The true model $\bar{\mathbf{w}}$ and feature matrices are generated by the same way as the linear regression experiment. The class label $y_i \in \{-1, 1\}$ is got by $y_i = \text{sign}(X_i, \bar{\mathbf{w}})$. Besides the selection recall, we also compare the classification error. In Figure 3, we can see that the superiority of our method is even more significant in the classification experiment. Although the overall sparsity has the lowest selection recall, it still has a similar classification error with the methods that consider row or column groups.




Figure 2. Selection recall and successful recovery rate for least square loss.




Figure 3. Selection recall and classification error for squared hinge loss.

5.2. Application in Crowdsourcing

This section applies the proposed method to the worker-task assignment problem in crowdsourcing. Take the im-

age labeling task as an example. Given n workers and m images, each image can be assigned to multiple workers and each worker can label multiple images. The predicted label for each image is decided by all the labels provided by the assigned workers and the quality of each worker on the image. The goal is to maximize the expected prediction accuracy based on the assignment. Let $X \in \{0, 1\}^{n \times m}$ be the assignment matrix, i.e. $X_{ij} = 1$ if assign the i -th worker to j -th task, otherwise $X_{ij} = 0$. $Q \in [0, 1]^{n \times m}$ is the corresponding quality matrix, which is usually estimated from the golden standard test (Ho et al., 2013). The whole formulation is defined to maximize the average of the expected prediction accuracy over m tasks over a TVCS constraint:

$$\begin{aligned} \max_X \quad & \frac{1}{m} \sum_{j=1}^m \mathcal{E}_{\text{acc}}(Q_{\cdot,j}, X_{\cdot,j}) \\ \text{subject to} \quad & \sum_{i=1}^n X_{ij} \leq s^{\text{worker}}, \forall j; \\ & \sum_{j=1}^m X_{ij} \leq s^{\text{task}}, \forall i; \\ & \sum_{i,j} X_{ij} \leq s^{\text{total}}, X \in \{0, 1\}^{n \times m} \end{aligned} \quad (8)$$

where $\mathcal{E}_{\text{acc}}(\cdot, \cdot)$ is the expected prediction accuracy, s^{worker} is the ‘‘worker sparsity’’, i.e. the largest number of assigned workers for each task, and s^{task} is the ‘‘task sparsity’’, i.e. each worker can be assigned with at most s^{task} tasks, and s^{total} is the total sparsity to control the budget, i.e., the maximal number of assignment. In image labeling task, we assume that each image can only have two possible classes and the percentage of images in each class is one half. We use the Bayesian rule to infer the predicted labels given the workers’ answer. Here we consider the binary classification task. Let $\mathbf{y}_j \in \{1, 0\}$ be the true label of the j -th task and $\hat{\mathbf{y}}_j$ be the prediction given labels by selected workers, i.e.,

$$\hat{\mathbf{y}}_j = \begin{cases} 0, & \text{if } \mathbb{P}(\mathbf{y}_j = 1 | \hat{Y}_{\Omega_j, j}) < \mathbb{P}(\mathbf{y}_j = 0 | \hat{Y}_{\Omega_j, j}); \\ 1, & \text{otherwise} \end{cases}$$

where \hat{Y}_{ij} is the i -th worker’s prediction on j -th task. Set Ω_j contains the indices of the selected workers for j -th task, i.e. $X_{ij} = 1, \forall i \in \Omega_j$, and $X_{i'j} = 0, \forall i' \notin \Omega_j$

Then $\mathcal{E}_{\text{acc}}(Q_{\cdot,j}, X_{\cdot,j})$ can be defined in the following:

$$\mathcal{E}_{\text{acc}}(Q_{\cdot,j}, X_{\cdot,j}) = \mathbb{P}(\hat{\mathbf{y}}_j = 1, \mathbf{y}_j = 1) + \mathbb{P}(\hat{\mathbf{y}}_j = 0, \mathbf{y}_j = 0)$$

By this way, the expected accuracy will not be continuous, so we use smooth function to approximate the expected accuracy and adopt the stochastic gradient with the proposed

projection operator to optimize it. Due to the space limitation, the detailed derivation of the objective formulation can be found in the supplemental material.

We conduct experiment for crowdsourcing task assignment on synthetic data. Specifically, we generate the quality matrix Q from uniformly random distribution with interval $[0.5, 0.9]$. The prior probability $\mathbb{P}(\mathbf{y}_j = 1)$ and $\mathbb{P}(\mathbf{y}_j = 0)$ are set as 0.5 for all the tasks.

To avoid evaluating the expectation term, we apply the stochastic iterative hard thresholding framework (Nguyen et al., 2014). Each iteration we get $\hat{Y}|\mathbf{y}_j = 1$ and $\hat{Y}|\mathbf{y}_j = 0$ by sampling based on Q i.e. $\mathbb{P}(\hat{Y}_{ij} = 1|\mathbf{y}_j = 1) = Q_{ij}$, $\mathbb{P}(\hat{Y}_{ij} = 0|\mathbf{y}_j = 0) = Q_{ij}$. Then we can get a stochastic gradient based on the sampled \hat{Y} .

Besides the proposed formulation (8), we evaluate the random assignment algorithm and the Q-based linear programming (Ho et al., 2013). The random assignment algorithm widely used in practice is the most straightforward approach: given the total assignment budget s^{total} and the restrictions (s^{worker} and s^{task}) for workers and tasks, randomly assign tasks to the workers. The Q-based linear programming uses the linear combination of Q_{ij} over i to evaluate the overall accuracy on task j for simpler formulation. In addition, it does not consider the restriction on tasks, thus it may assign lots of workers to a difficult task². To make a fair comparison, the task restriction is added into this method. To get the assignment result which satisfies the task and worker restriction, we use our projection operator in the other methods too.

We evaluate the experiments on different value of $s^{\text{task}}, s^{\text{worker}}$ by setting them as different ratios of the total number of tasks and workers. The overall sparsity is set by the same way as in Section 5.1. To measure the performance, we compare the sampled expected accuracy. The samples (i.e., \hat{Y}) are independent to the samples used in training. Figure 4 shows the comparison of the expected accuracy of the three approaches. We can observe that the accuracy increases with larger ratio (i.e. more assignments). The random assignment strategy needs more assignments to get the same accuracy compared with the other two methods.

5.3. Application in Identification of Gene Regulatory Networks

In this section, we apply the projection operator to the identification of gene regulatory networks (GRN).

²A “difficult” task means that all workers’ qualities are low on this task.




Figure 4. Expected accuracy of crowdsourced classification.

Background. Gene regulatory network represents the relations between different genes, which plays important roles in biological processes and activities by controlling the expression level of RNAs. There is a well-known biological competition named DREAM challenge about identifying GRN. Based on the time series gene expression data which are RNAs’ level along time sequence, contestants are required to recover the whole gene network of given size. One popular way to infer GRN is to utilize the sparse property of GRN: e.g., one gene in the network is only related to a small number of genes and we already know that there exists no relationship between some genes. Therefore, the amount of edges connecting to one vertex is far less than the dimension of the graph. It is a practical case of row-wise and column-wise sparsity for matrix. We could apply the projection operator to constrain the number of edges related to each vertex to identify the whole network. Recently, the dynamic Bayesian network (DBN) (Zou & Conzen, 2005) is supposed to be an effective model to recover GRNs. The RNAs’ level of all genes in GRN at time t is stored in gene expression vector $\mathbf{x}_t \in \mathbb{R}^N$, where each entry corresponds to one gene respectively and N is the number of genes in GRN. Hence, We define the total amount of time points in the experiment as T . Gene activity model is usually assumed to be

$$\mathbf{x}_{t+1} = P\mathbf{x}_t + \mathbf{e}_t, \quad t = 1 \dots T - 1,$$

where $P \in \mathbb{R}^{N \times N}$ is the covariance matrix of GRN and $\mathbf{e}_t \in \mathbb{R}^N$ is Gaussian white noise. Then the difference of RNA levels between time points $t + 1$ and t , i.e. $\mathbf{y}_{t+1,t} \in \mathbb{R}^N$ is defined as follows:

$$\mathbf{y}_{t+1,t} := \mathbf{x}_{t+1} - \mathbf{x}_t = \bar{W}\mathbf{x}_t + \mathbf{e}_t, \quad t = 1 \dots T - 1,$$

where $\bar{W} = P - I$ is the true sparse N -by- N matrix. Therefore, the GRN is only considered between different genes and we eliminate edges whose start and end vertex are the same. We define that $Y := [\mathbf{y}_{2,1}, \dots, \mathbf{y}_{T,T-1}] \in \mathbb{R}^{N \times (T-1)}$ and $X := [\mathbf{x}_1, \dots, \mathbf{x}_{T-1}] \in \mathbb{R}^{N \times (T-1)}$. The objective function is

$$f(W) = \frac{1}{2} \|Y - \bar{W}X\|_{\mathcal{F}}^2 = \frac{1}{2} \sum_{t=1}^{T-1} \|(\mathbf{x}_{t+1} - \mathbf{x}_t) - \bar{W}\mathbf{x}_t\|^2.$$

On The Projection Operator to A Three-view Cardinality Constrained Set

	SN	SP	ACC	F-measure	MCC	AUC
Our Method	0.6875±0.0295	0.7397±0.0319	0.7119±0.0305	0.7126±0.0306	0.4264±0.0611	0.7136±0.0306
GENIE3	0.5611±0.0277	0.4984±0.0547	0.5319±0.0244	0.5279±0.0277	0.0595±0.0547	0.5662±0.0244
CLR	0.5167±0.0583	0.4476±0.1147	0.4844±0.0575	0.4795±0.0583	-0.0357±0.1147	0.5210±0.0575
TIGRESS	0.1333±0.0541	0.8302±0.0367	0.4585±0.0374	0.2258±0.0817	-0.0552±0.1061	0.5567±0.0358
PCC	0.5042±0.0124	0.4333±0.0245	0.4711±0.0101	0.4661±0.0124	-0.0625±0.0245	0.5091±0.0101
ARACNE	0.1167±0.0519	0.9127±0.0579	0.4881±0.0197	0.2051±0.0519	0.0479±0.0579	0.5808±0.0197
MINET	0.5764±0.0425	0.5381±0.0888	0.5585±0.0458	0.5547±0.0425	0.1147±0.0888	0.5910±0.0458

Table 1. Performance evaluation of our method and six other state-of-art methods

Time-course Gene Expression Data. To evaluate our method, we employ GeneNetWeaver (Marbach et al., 2009; Schaffter et al., 2011), the official DREAM Challenge tool for time-series expression data generation. With typical gene network structure and ordinary differential equation (ODE) models, GeneNetWeaver will produce the time-course gene expression data at pre-specified time points. In the simulation studies, we control the size of gene network to be $N = 30$ vertexes and the gene expression data are generated under 10% Gaussian white noise.

The network is shown in Figure 5. In this Figure, it is clear that one gene only has a few connections to other genes. Therefore, the GRN is sparse and we are able to restrict the in-degree and out-degree of every vertex by representing the network as a matrix and controlling the sparsity within each row and column.

Performance evaluation. Six commonly-used criteria are considered to measure the performance, i.e., sensitivity (SN), specificity (SP), accuracy (ACC), F-measure, Matthews correlation coefficient (MCC), and the Area Under ROC Curve (AUC):

$$\begin{aligned}
 \text{SN} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, & \text{SP} &= \frac{\text{TN}}{\text{TN} + \text{FP}}, \\
 \text{ACC} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}}, \\
 \text{F-measure} &= \frac{2 \times \text{SN} \times \text{SP}}{\text{SN} + \text{SP}}, \\
 \text{MCC} &= \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}},
 \end{aligned}$$

where TP and TN denote the true positive and true negative, and FP and FN denote the false positive and false negative, respectively. With these criteria, we compare the performance of our method with six representative algorithms, including PCC, ARACNE (Margolin et al., 2006), CLR (Faith et al., 2007), MINET (Meyer et al., 2008), GENIE3 (Huynh-Thu et al., 2010), TIGRESS (Hauray et al., 2012). The results are summarized in Table 1. Our method outperforms other six state-of-art methods: the AUC of our method achieve 0.7 higher which is far more than other methods; 5 out of 6 different measure show that our method has significant advantage compared to other algorithms.




Figure 5. This gene regulatory network contains 30 vertexes which are represented by blue circles. One edge starts at gene 1 and ends at gene 2 if gene 1 has influence on gene 2.

6. Conclusion

This paper considers the TVCS constrained optimization, motivated by the intrinsic restrictions for many important applications, for example, in bioinformatics, recommendation system, and crowdsourcing. To solve the cardinality constrained problem, the key step is the projection onto the cardinality constraints. Although the projection onto the overlapped cardinality constraints is NP-hard in general, we prove that if the TVCS condition is satisfied the projection can be reduced to a linear programming. We further prove that there is an iterative algorithm which finds an integer solution to the linear programming within time complexity $O((p + |\mathcal{G}|) \log_{\alpha} \frac{1}{R})$, where R is the distance from the initial point to the optimization solution and $\alpha < 1$ is the convergence rate. We finally use synthetic experiments and two interesting applications in bioinformatics and crowdsourcing to validate the proposed TVCS model.

Acknowledgements

This project is supported in part by the NSF grant CNS-1548078 and the NEC fellowship.

References

- Bach, Francis, Jenatton, Rodolphe, Mairal, Julien, Obozinski, Guillaume, et al. Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468, 2012.
- Baraniuk, Richard G, Cevher, Volkan, Duarte, Marco F, and Hegde, Chinmay. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- Candes, Emmanuel J, Romberg, Justin K, and Tao, Terence. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.
- Cevher, Volkan, Indyk, Piotr, Hegde, Chinmay, and Baraniuk, Richard G. Recovery of clustered sparse signals from compressive measurements. Technical report, DTIC Document, 2009.
- El Halabi, Marwa and Cevher, Volkan. A totally unimodular view of structured sparsity. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pp. 223–231, 2015.
- Faith, Jeremiah J, Hayete, Boris, Thaden, Joshua T, Mogno, Iliaria, Wierzbowski, Jamey, Cottarel, Guillaume, Kasif, Simon, Collins, James J, and Gardner, Timothy S. Large-scale mapping and validation of *Escherichia coli* transcriptional regulation from a compendium of expression profiles. *PLoS Biol*, 5(1):e8, 2007.
- Hauray, Anne-Claire, Mordelet, Fantine, Vera-Licona, Paola, and Vert, Jean-Philippe. Tigress: trustful inference of gene regulation using stability selection. *BMC systems biology*, 6(1):145, 2012.
- Hegde, Chinmay, Duarte, Marco F, and Cevher, Volkan. Compressive sensing recovery of spike trains using a structured sparsity model. In *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- Hegde, Chinmay, Indyk, Piotr, and Schmidt, Ludwig. Approximation algorithms for model-based compressive sensing. *Information Theory, IEEE Transactions on*, 61(9):5129–5147, 2015a.
- Hegde, Chinmay, Indyk, Piotr, and Schmidt, Ludwig. A nearly-linear time framework for graph-structured sparsity. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 928–937, 2015b.
- Ho, Chien-Ju, Jabbari, Shahin, and Vaughan, Jennifer Wortman. Adaptive task assignment for crowd-sourced classification. In *Proceedings of The 30th International Conference on Machine Learning*, pp. 534–542, 2013.
- Hoffman, Alan J. On approximate solutions of systems of linear inequalities. In *Selected Papers Of Alan J Hoffman: With Commentary*, pp. 174–176. 2003.
- Huynh-Thu, Vn Anh, Irrthum, Alexandre, Wehenkel, Louis, and Geurts, Pierre. Inferring regulatory networks from expression data using tree-based methods. *PLoS one*, 5(9):e12776, 2010.
- Jenatton, Rodolphe, Mairal, Julien, Obozinski, Guillaume, and Bach, Francis. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12(Jul):2297–2334, 2011.
- Kong, Deguang, Fujimaki, Ryohei, Liu, Ji, Nie, Feiping, and Ding, Chris. Exclusive feature learning on arbitrary structures via $l_{1,2}$ -norm. In *Advances in Neural Information Processing Systems*, pp. 1655–1663, 2014.
- Liu, Ji and Wright, Stephen J. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.
- Marbach, Daniel, Schaffter, Thomas, Mattiussi, Claudio, and Floreano, Dario. Generating realistic in silico gene networks for performance assessment of reverse engineering methods. *Journal of computational biology*, 16(2):229–239, 2009.
- Margolin, Adam A, Nemenman, Ilya, Basso, Katia, Wiggins, Chris, Stolovitzky, Gustavo, Favera, Riccardo D, and Califano, Andrea. Aracne: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC bioinformatics*, 7(Suppl 1):S7, 2006.
- Meyer, Patrick E, Lafitte, Frederic, and Bontempi, Gianluca. minet: Ar/bioconductor package for inferring large transcriptional networks using mutual information. *BMC bioinformatics*, 9(1):461, 2008.
- Needell, Deanna and Tropp, Joel A. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- Nguyen, Nam, Chin, Sang, and Tran, Trac D. A unified iterative greedy algorithm for sparsity constrained optimization. 2012.

- Nguyen, Nam, Needell, Deanna, and Woolf, Tina. Linear convergence of stochastic iterative greedy algorithms with sparse constraints. *arXiv preprint arXiv:1407.0088*, 2014.
- Olshausen, Bruno A and Field, David J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.
- Papadimitriou, Christos H and Steiglitz, Kenneth. *Combinatorial optimization: algorithms and complexity*. Courier Corporation, 1982.
- Schaffter, Thomas, Marbach, Daniel, and Floreano, Dario. Genenetweaver: in silico benchmark generation and performance profiling of network inference methods. *Bioinformatics*, 27(16):2263–2270, 2011.
- Yang, Haichuan, Huang, Yijun, Tran, Lam, Liu, Ji, and Huang, Shuai. On benefits of selection diversity via bilevel exclusive sparsity. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*. IEEE, 2016.
- Yuan, Ming and Lin, Yi. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- Yuan, Xiaotong, Li, Ping, and Zhang, Tong. Gradient hard thresholding pursuit for sparsity-constrained optimization. In *Proceedings of The 31st International Conference on Machine Learning*, pp. 127–135, 2014.
- Yuan, Xiaotong, Li, Ping, and Zhang, Tong. Exact recovery of hard thresholding pursuit. In *Advances in Neural Information Processing Systems*, pp. 3558–3566, 2016.
- Zhang, Tong. On the consistency of feature selection using greedy least squares regression. *Journal of Machine Learning Research*, 10(Mar):555–568, 2009.
- Zhou, Yang, Jin, Rong, and Hoi, Steven CH. Exclusive lasso for multi-task feature selection. In *AISTATS*, volume 9, pp. 988–995, 2010.
- Zou, Min and Conzen, Suzanne D. A new dynamic bayesian network (dbn) approach for identifying gene regulatory networks from time course microarray data. *Bioinformatics*, 21(1):71–79, 2005.