

Algorithms for Generalized Cluster-wise Linear Regression

Young Woong Park* Yan Jiang[†] Diego Klabjan[‡] Loren Williams[§]

Apr 3, 2015

Abstract

Cluster-wise linear regression (CLR), a clustering problem intertwined with regression, is to find clusters of entities such that the overall sum of squared errors from regressions performed over these clusters is minimized. We generalize the CLR problem by allowing each entity to have more than one observation, and refer to it as generalized CLR. We propose an exact mathematical programming based approach relying on column generation, a column generation based heuristic algorithm that clusters predefined groups of entities, a metaheuristic genetic algorithm with adapted Lloyd’s algorithm for K-means clustering, and a two-stage approach that performs clustering first and regression second for solving generalized CLR. We examine the performance of our algorithms on a stock keeping unit (SKU) clustering problem employed in forecasting halo and cannibalization effects in promotions using real-world retail data from a large supermarket chain. In the SKU clustering problem, the retailer needs to cluster SKUs based on their seasonal effects in response to promotions. The seasonal effects are the results of regressions with predictors being promotion mechanisms and seasonal dummies performed over clusters generated. We compare the proposed algorithms for the SKU problem with real-world and synthetic data.

1 Introduction

Clustering is a commonly encountered problem in many areas such as marketing, engineering, and biology, among others. In a typical clustering problem, the goal is to group entities together according to a certain similarity measure. Such a measure can be defined in many different ways, and it determines the complexity of solving the relevant clustering problem. Clustering problems with the similarity measure defined by regression errors is especially challenging because it is coupled with regression.

Consider a retailer that needs to forecast sales at the stock keeping unit (SKU) level for different promotional plans and mechanisms (e.g., 30% off the selling price) using a linear regression model. A SKU is a unique identifying number that refers to a specific item in inventory. Each SKU is often used to identify product, product size, product type, and the manufacturer. Seasonality is an important predictor and is modeled using an indicator dummy input variable for each season, with the length of one season being one week. The usable data for each SKU is limited comparing to the possible number of parameters to estimate, among which the seasonality dummies compose a large proportion. More significant and useful statistical results can be obtained by clustering SKUs with similar seasonal effects from promotions together, and estimating seasonality dummies for a cluster instead of a single SKU. However, the seasonal effects of SKUs correspond to regression coefficients, which can only be obtained after grouping SKUs with similar seasonality.

A two-stage method can be used to solve such difficult clustering problems that are intertwined with regression. In the first stage, entities are clustered based on certain approximate measures of their regression

*ywpark@u.northwestern.edu, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, USA.

[†]jiangyan1984@gmail.com, Sears Holdings Corporation, Hoffman Estates, IL

[‡]d-klabjan@northwestern.edu, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, USA.

[§]loren.williams@predictix.com, Predictix LLC, Atlanta, GA

coefficients. In the second stage, regressions are performed over the resultant clusters to obtain estimates for the regression coefficients for each cluster. However, good approximate measures are difficult to obtain a priori before carrying out the regressions. A better alternative is to perform clustering and regression simultaneously, which can be achieved through cluster-wise linear regression (CLR), which is also referred to as “regression clustering” in the literature. Other application areas of CLR include marketing, pavement condition prediction, and spatial modeling and analysis. More details about these other application areas can be found in Openshaw [20], DeSarbo and Cron [7], DeSarbo [6], and Luo and Chou [17].

The CLR problem bears connection to the minimum sum-of-squares clustering (MSSC) problem, the objective of which is to find clusters that minimize the sum of squared distances from each entity to the centroid of the cluster which it belongs to. Contrary to clustering entities directly based on distances, CLR generates clusters according to the effects that some independent variables have on the response variable of a preset regression model. Each entity is represented by a set of observations of a response variable and the associated predictors. CLR is to group entities with similar regression effects into a given number of clusters such that the overall sum of squared residuals within clusters is minimal. Although the MSSC problem has been extensively studied by researchers from various fields (e.g., statistics, optimization, and data mining), the work for the CLR problem is limited, most of which concerns adapting the Lloyd’s algorithm based heuristic algorithms of the MSSC problem to the CLR problem. The Lloyd’s algorithm starts randomly from some initial partition of clusters, then calculates the centroids of clusters, and assigns entities to their closest centroids until converging to a local minimum. To the best of our knowledge, there is no algorithm solving CLR exactly, and thus the performance of these heuristic algorithms is unknown.

We tackle the problem of clustering entities based on their regression coefficients by modeling it as a generalized CLR problem, in which we allow each entity to have more than one observation. We propose both a mixed integer quadratic program formulation and a set partitioning formulation for generalized CLR. Our mixed integer quadratic program formulation is more general than the one proposed by Bertsimas and Shioda [4], which cannot be directly applied to the SKU clustering problem since they assume each clustering entity to have only one observation and this assumption does not hold for the SKU clustering problem. We identify a connection between the generalized CLR and MSSC problems, through which we prove NP-hardness of the generalized CLR problem. Column generation is an algorithmic framework for solving large-scale linear and integer programs. Vanderbeck and Wolsey [22] and Barnhart et al. [3] overview column generation for solving large integer program. We design a column generation (CG) algorithm for the generalized CLR problem using its set partitioning formulation. The corresponding pricing problem is a mixed integer quadratic program, which we show to be NP-hard. To handle larger instances in the column generation framework, we also propose a heuristic algorithm, referred to as the CG Heuristic algorithm. This heuristic algorithm, inspired by Bertsimas and Shioda [4], first clusters entities to a small number of groups and then performs our column generation algorithm on these groups of entities. In addition, we propose a metaheuristic algorithm, named the GA-Lloyd algorithm, which uses an adapted Lloyd’s clustering algorithm to find locally optimal partitions and relies on the genetic algorithm (GA) to escape local optimums. Furthermore, we introduce a two-stage approach, used frequently in practice due to its simplicity, which performs clustering first and regression second. We test our algorithms using real-world data from a large retail chain. We compare the performance of the GA-Lloyd, the CG Heuristic, and the two-stage algorithms on two larger instances with 66 and 337 SKUs, corresponding to two representative subcategories under the retailer’s product hierarchy. We observe that the GA-Lloyd algorithm performs much better than the two-stage algorithm. The CG Heuristic algorithm is able to produce slightly better results than the GA-Lloyd algorithm for smaller instances, but at the cost of much longer running time. The GA-Lloyd algorithm performs the best and identifies distinctive and meaningful seasonal patterns for the tested subcategories. In addition, we find that the column generation algorithm is able to solve the SKU clustering problem with at most 20 SKUs to optimality within reasonable computation time. We benchmark the performance of the GA-Lloyd and CG Heuristic algorithms against the optimal solutions obtained by the column generation algorithm to find that both algorithms obtain close to optimal solutions.

The contributions of our work are as follows.

1. We are the first to model and solve the SKU clustering problem, commonly encountered in retail predictive modeling, through generalized CLR.
2. We propose three heuristic algorithms for the generalized CLR problem, including the CG Heuristic algorithm, the GA-Lloyd algorithm, and the two-stage approach.

3. We propose an exact column generation algorithm that enables us to evaluate the performance of the heuristic algorithms.
4. We prove NP-hardness of the generalized CLR problem and NP-completeness of the pricing problem of the column generation algorithm.

Note that the number of clusters is a parameter in the generalized CLR problem that needs to be decided by user beforehand or by enumeration. Although we provide comparison of models with different number of clusters for real-world data in Section 4.2, it is not straightforward to develop a universal rule for deciding the number of clusters. This is also a hard task for MSSC and CLR. Hence, in this paper, we assume that the target number of clusters is given in advance.

The rest of the paper is organized as follows. In Section 2, we introduce both the mixed integer quadratic program and the set partitioning formulations of the generalized CLR problem. We draw the connection between the generalized CLR and MSSC problems, and prove NP-hardness of the former through this connection. In Section 3, we present the exact column generation algorithm, the CG Heuristic algorithm, the GA-Lloyd heuristic algorithm, and the two-stage algorithm. The pricing problem of the column generation algorithm is shown to be NP-complete. In Section 4, we present numerical experiments to test the performance of all proposed algorithms. The literature review is discussed next.

1.1 Literature Review

Due to the novelty of this work, the scientific literature revealed that no previous work has been conducted that comprehensively tackles the generalized CLR problem. However, an extensive collection has been proposed for the typical CLR problem, which can potentially be adapted to tackle the generalized CLR problem.

The algorithms proposed for the typical CLR problem are mainly heuristics bearing close similarity to the algorithms for the MSSC problem. For example, Spath [21] proposes an exchange algorithm which, starting from some initial clusters, exchanges two items between two clusters if a cost reduction is observed in the objective function. DeSarbo [6] presents a simulated annealing method to escape local minimums. Muruzabal et al. [19] used a self organizing map to perform clusterwise regression.

On mathematical programming-based heuristics, Lau et al. [16] propose a nonlinear programming formulation that it is solved approximately using commercial solvers with no guarantee to find a global optimum. Their algorithm’s performance depends heavily on the initial clusters. This initial-cluster dependency is overcome by the K-harmonic means clustering algorithm proposed by Zhang [24]. Moreover, Bertsimas and Shioda [4] introduce a compact mixed-integer linear formulation for a slight variation of the CLR problem with the sum of the absolute error as the objective. Their algorithm first divides entities into a small number of clusters, and then feeds these clusters into their mixed integer program.

There are three key differences between these works and the one we propose in this paper. First, we provide both a quadratic mixed-integer program formulation and a set partition formulation of the generalized CLR problem. The former is a generalization of the formulation in Bertsimas and Shioda [4], and the latter has never been proposed for the CLR problem. Second, we propose two new heuristics, namely the CG Heuristic algorithm and the GA-Lloyd algorithm for the generalized CLR problem. Third, unlike the aforementioned works, we provide a measure of our algorithms’ performance in terms of the optimality gap with respect to optimal solutions obtained through our exact method.

There is another stream of research for the CLR problem that assumes a distribution function for regression errors where each entity is assigned to each cluster with a certain probability, i.e., using “soft” assignments. For example, DeSarbo and Cron [7] propose a finite conditional mixture maximum likelihood methodology, which assumes normal distribution for regression errors and is solved through the expectation maximization algorithm. Since then, a large number of mixture regression models have been developed, including probit and logit mixture regression models as examples. Lau et al. [16] compare the performance of the expectation maximization algorithms with their nonlinear programming-based algorithm. D’Urso et al. [10] proposed to integrate fuzzy clustering and fuzzy regression. A recent work of Ingrassia et al [14] uses linear t cluster-weighted models for clustering regression. These model-based approaches allow residual variances to differ between clusters, which the least squares approaches do not allow. In the soft assignment setting, an entity can be assigned to the cluster of highest probability. We restrict the scope of our review and comparison to least squares approaches because the objective functions are different. The reader is referred to Wedel and DeSarbo [23] and Hennig [13] for reviews.

The algorithms for the MSSC problem are instructive to solving the CLR problem. There are abundant papers for solving the MSSC problem. Hansen and Jaumard [12] survey various forms of clustering problems and their solution methods, including MSSC, from a mathematical programming point of view. In their survey, solution methods for the MSSC problem include dynamic programming, branch-and-bound, cutting planes, and column generation methods. All these algorithms do not scale well to large size instances or in higher dimensional spaces. Heuristics are also considered, including Lloyd’s like algorithms (e.g., K-Means and H-Means) and metaheuristics such as simulated annealing, tabu search, genetic algorithms and variable neighborhood search. With respect to mathematical programming approaches, du Merle et al. [8] propose an interior point algorithm to exactly solve the MSSC problem. Aloise et al. [2] improve the algorithm of du Merle et al. [8] by exploiting the geometric characteristics of clusters, which enables them to solve much larger instances.

2 Problem Formulations

2.1 Mixed Integer Quadratic Program Formulation

We first provide a mixed integer quadratic formulation for the generalized CLR problem. This formulation reveals a close connection between the generalized CLR and MSSC problems, which enables us to show that the generalized CLR problem is NP-hard.

Consider set $\{1, 2, \dots, I\}$ of I entities. Each entity $i \in I$ has L observations of dependent variable $\mathbf{y}_i = (y_{i1}, y_{i2}, \dots, y_{iL})$, and J independent variables $\mathbf{x}_{i1}, \mathbf{x}_{i2}, \dots, \mathbf{x}_{iJ}$ with $\mathbf{x}_{ij} = (x_{ij1}, x_{ij2}, \dots, x_{ijL})$ for any $j \in [J]$. In practice the number of entities L depends on i , but we do not show this dependency for improved readability. (For each integer g we introduce $[g] = \{1, \dots, g\}$.) Observation y_{il} is associated with independent variables $x_{i1l}, x_{i2l}, \dots, x_{iJl}$. Note that vectors are represented in bold symbols. We want to divide these I entities into a partition \mathbf{C} of K clusters where $\mathbf{C} = (C_1, C_2, \dots, C_K)$, $C_i \cap C_j = \emptyset$ for any $i \neq j$, and $\bigcup_{k \in [K]} C_k = [I]$. The minimum size of a cluster is n , which implies $|C_k| \geq n$ for any $k \in [K]$ where $|C_k|$ denotes the cardinality of cluster C_k . Note that the number of observations pertaining to a cluster is at least nL . The minimum size constraints are imposed to ensure that there are enough observations for each cluster so that the regression over each cluster generates statistically significant results. Further, in order to avoid collinearity in the regression models, we require $L \cdot n > J + 1$. We also require $I \geq K \cdot n$ such that there is always a feasible solution. The generalized CLR problem is formulated as follows:

$$\min \sum_{i=1}^I \sum_{l=1}^L t_{il}^2 \quad (1)$$

$$t_{il} - (y_{il} - \sum_{j=1}^J \beta_{kj} x_{ijl}) + M(1 - z_{ik}) \geq 0 \quad i \in [I], k \in [K], l \in [L] \quad (2)$$

$$t_{il} + (y_{il} - \sum_{j=1}^J \beta_{kj} x_{ijl}) + M(1 - z_{ik}) \geq 0 \quad i \in [I], k \in [K], l \in [L] \quad (3)$$

$$\sum_{k=1}^K z_{ik} = 1 \quad i \in [I] \quad (4)$$

$$\sum_{i=1}^I z_{ik} \geq n \quad k \in [K] \quad (5)$$

$$z_{ik} \in \{0, 1\} \quad i \in [I], k \in [K]$$

$$t_{il} \geq 0 \quad i \in [I], l \in [L]$$

$$\beta_{kj} \text{ unconstrained} \quad k \in [K], j \in [J],$$

where z_{ik} is a binary variable, which is equal to one if and only if entity i is assigned to cluster C_k . Value M , referred to as big M in the optimization literature, is a large positive constant. Due to constraints (2)

and (3), t_{il} is equal to the absolute error for the corresponding observation y_{il} in the optimal solution, and $\beta_k = (\beta_{k1}, \beta_{k2}, \dots, \beta_{kJ})$ are the regression coefficients for cluster C_k , which are decision variables. The role of M is to enforce constraints (2) and (3) only when they are needed (entity i is assigned to cluster k). In detail, if $z_{ik} = 1$, then we have $t_{il} - (y_{il} - \sum_{j=1}^J \beta_{kj} x_{ijl}) \geq 0$, and $t_{il} + (y_{il} - \sum_{j=1}^J \beta_{kj} x_{ijl}) \geq 0$, which implies $t_{il} = |y_{il} - \sum_{j=1}^J \beta_{kj} x_{ijl}|$ because we are minimizing the sum of t_{il}^2 . If $z_{ik} = 0$, constraints (2) and (3) require t_{il} to be greater than a negative number, which holds trivially due to the existence of the nonnegativity constraint on t_{il} . Constraint (4) requires that every entity is assigned to one cluster, and (5) imposes the limit on the cardinality of each cluster.

Unlike the CLR problem, the generalized CLR allows each entity to have more than one observation, which implies that L can be greater than one. The mixed integer linear program formulation for the CLR problem in Bertsimas and Shioda [4] has L equal to one, and does not have the cluster cardinality constraint (5). Besides, their objective function is the sum of the absolute errors while ours is the sum of squared errors.

Our SKU clustering problem based on the seasonal effects can be modeled as the generalized CLR problem. The entities to cluster are SKUs. The response variable \mathbf{y}_i corresponds to a vector of weekly sales for SKU i . The independent variables \mathbf{x}_i 's include promotional predictors such as promotion mechanisms, percentage discount, and seasonal dummies for SKU i .

Aloise et al. [1] showed NP-hardness of the MSSC problem in a general dimension when the number of clusters is two. General dimension means that the size of the vectors to be clustered is not a constant but part of the input data. A similar statement can be made for the generalized CLR problem with the proof available in Appendix A.

Theorem 1 *The generalized CLR problem with two clusters in a general dimension is NP-hard.*

With the formulation presented by (1)–(5), we can solve the generalized CLR problem using any commercial optimization software that can handle quadratic mixed integer programs. However, this formulation suffers from two drawbacks, which makes it intractable for large instances. The first one relates to big M . Optimality of the solution and efficiency of integer programming solvers depend on a tight value of M . However, a tight M cannot be estimated easily. The second one involves the symmetry of feasible solutions. Any permutation of clusters yields the same solution, yet it corresponds to different decision variables. Symmetry unnecessarily increases the search space, and renders the solution process inefficient. To overcome the symmetry problem, we propose a set partitioning formulation.

2.2 Set Partitioning Formulation

Let \mathcal{S} denote the set of all clusters of entities with the cardinality equal to or greater than n , i.e., $\mathcal{S} = \{S \subseteq [I], |S| \geq n\}$. Let a_{iS} equal to one if entity i belongs to cluster S , and equal to zero otherwise. Let c_S denote the cost of cluster S , which is equal to the sum of squared errors when performing the regression over cluster S . Introducing binary variables

$$z_S = \begin{cases} 1 & \text{if cluster } S \text{ is selected,} \\ 0 & \text{otherwise,} \end{cases}$$

the generalized CLR problem can be formulated as:

$$\min \sum_{S \in \mathcal{S}} c_S z_S \tag{6}$$

$$\sum_{S \in \mathcal{S}} z_S = K \tag{7}$$

$$\sum_{S \in \mathcal{S}} a_{iS} z_S = 1 \quad i \in [I] \tag{8}$$

$$z_S \in \{0, 1\} \quad S \in \mathcal{S}.$$

Constraint (7) ensures that the number of clusters in the partition is K and constraint (8) guarantees that each entity occurs in only one cluster within the partition.

3 Algorithms

3.1 Column Generation (CG) Algorithm

The set partitioning formulation has an exponential number of binary variables. It is very challenging to solve even its linear programming relaxation because there are so many decision variables. To solve large-scale linear and integer programs, column generation algorithms have been used in the literature. The reader is referred to Vanderbeck and Wolsey [22] and Barnhart et al. [3] for reviews of column generation for solving large-scale integer programs. In our work, we employ column generation to handle its linear programming relaxation. At the high level, column generation can be understood as iteratively expanding set \mathcal{S} (a subset of \mathcal{I}) in (6) - (8) by adding attractive candidate cluster S to \mathcal{S} . The key challenge is how to select S . The word column is used because adding cluster S to \mathcal{S} is equivalent to adding a column in the matrix form of (6) - (8).

The column generation algorithm, referred to as the CG algorithm, starts by solving the restricted master problem which has the same formulation as the master problem (6)-(8), but with set \mathcal{S} replaced by \mathcal{S} , a smaller subset of columns. Recall that a column represents a cluster (subset of entities $[I]$). We start the algorithm with small candidate clusters rather than \mathcal{S} , the set of all possible subsets of $[I]$. The algorithmic framework is presented in Algorithm 1, which follows the general column generation scheme. In Line 1, the initial subset of columns in \mathcal{S} are randomly generated. In Line 3, optimal dual variables are obtained by solving the restricted master problem and then serve as input to the pricing problem, which will be introduced hereafter, to calculate the smallest reduced cost column. In Line 4, the pricing problem returns a column with the smallest reduced cost. In Lines 5-10, if the reduced cost is nonnegative, then we conclude that the master problem is solved optimally. Otherwise, we add the column with the smallest reduced cost to the restricted master problem and repeat the process.

Algorithm 1 CG

```

1: Randomly generate  $\mathcal{S}$  (a small subset of  $\mathcal{I}$ )
2: while not optimal do
3:   Solve master problem (6) – (8) and obtain dual solution
4:   Get a new cluster by solving pricing problem with input of dual solution from Line 3
5:   if the reduced cost is nonnegative then
6:     The algorithm is complete with the optimal partition of clusters
7:   else
8:     Add the cluster from Line 4 to the master problem
9:   end if
10: end while

```

The pricing problem

The pricing problem can be stated as follows. Let v be the dual variable for constraint (7), and π_i 's be the dual variables for constraint (8). The reduce cost for cluster S is $d_S = c_S - v - \sum_i \pi_i a_{iS}$, and thus the pricing problem reads:

$$\min_{|S| \geq n, \beta} \sum_{i \in S} \sum_{l=1}^L (y_{il} - \sum_{j=1}^J x_{ijl} \beta_j)^2 - \sum_{i \in S} \pi_i. \quad (9)$$

Note that we omit the subtraction of v in the formulation because it is a constant which does not change the optimal solution.

Theorem 2 *The pricing problem as stated in (9) is NP-complete.*

The proof is available in Appendix A. Introducing binary variables

$$z_i = \begin{cases} 1 & \text{if } i \in S, \\ 0 & \text{otherwise,} \end{cases}$$

the pricing problem can be formulated as a mixed integer quadratic program:

$$\min \sum_{i=1}^I \sum_{l=1}^L t_{il}^2 - \sum_{i=1}^I \pi_i z_i \quad (10)$$

$$t_{il} - (y_{il} - \sum_{j=1}^J \beta_j x_{ijl}) + N(1 - z_i) \geq 0 \quad i \in [I], l \in [L] \quad (11)$$

$$t_{il} + (y_{il} - \sum_{j=1}^J \beta_j x_{ijl}) + N(1 - z_i) \geq 0 \quad i \in [I], l \in [L] \quad (12)$$

$$\sum_{i=1}^I z_i \geq n \quad (13)$$

$$t_{il} \geq 0 \quad i \in [I], l \in [L]$$

$$z_i \in \{0, 1\} \quad i \in [I],$$

where N is a large positive constant. By using similar arguments as those for constraints (2) and (3), t_{il} is the absolute error for the corresponding observation y_{il} in the optimal solution if $i \in S$, and it is zero otherwise.

In the column generation algorithm, (10)–(13) are solved. Recall that reduced cost for cluster S is $d_S = c_S - v - \sum_i \pi_i a_{iS}$. It is easy to see that value $d_S + v$ is equivalent to the value of (10) with $z_i = 1$ for $i \in S$ and 0 otherwise. This follows from the fact that $d_S + v = c_S - \sum_{i \in S} \pi_i$ and c_S is modeled by variables t .

Column generation stabilization schemes

If the optimal solution obtained by CG is not integral, branching would have to be performed, i.e., a fractional variable z_S needs to be selected and two new problems created, the first one would impose $z_S = 0$ and the other one $z_S = 1$. However, the extensive evaluation conducted on Algorithm 2 revealed that no fractional solutions were provided by Algorithm 2. For this reason in the remainder we focus on column generation for solving the LP relaxation and not branching. Column generation is known to exhibit the tailing-off effect and for this reason we employ stabilized column generation of du Merle et al. [9].

The stabilized column generation algorithm for solving the CLR problem is illustrated in Algorithm 2. The algorithm takes input of stabilization parameters $\delta^{(0)}$ and $\xi^{(0)}$, and maximum allowed iterations k^{max} . In Line 1, we start with a set $\mathcal{S}^{(0)}$ of initial clusters of entities. For iteration k , in Line 3, we solve the stabilized master problem and get the optimal solution $(\mathbf{z}^{(k)}, \mathbf{q}^-, \mathbf{q}^+)$ and its corresponding dual solution $(\boldsymbol{\pi}^{(k)}, v)$, which provides input parameters for the pricing problem. The stabilized master problem additionally includes parameters $\delta^{(k)}$, $\xi^{(k)}$ and variables $\mathbf{q}^-, \mathbf{q}^+$ but is very similar to (6)–(8). See Appendix C for the actual formulation. By solving the pricing problem, we get a new cluster $S^{(k)}$ in Line 4. The reduced cost corresponding to this new cluster is equal to $c_S - \mathbf{a}_S^T \boldsymbol{\pi}^{(k)} - v$, where c_S is the sum of squared residuals when performing regression over this cluster, and $a_{iS} = 1$ if and only if $i \in S$. In Lines 5–6, if the reduced cost is nonnegative and \mathbf{q}^- and \mathbf{q}^+ are equal to zero, then the algorithm is complete with the optimal partition of clusters defined by $\mathbf{z}^{(k)}$. Otherwise, in Lines 8–12, we update \mathcal{S} and then if the reduced cost $c_S - \mathbf{a}_S^T \boldsymbol{\pi}^{(k)} - v$ is nonnegative, we update the stabilization parameters $\delta^{(k)}$, and $\xi^{(k)}$.

3.2 CG Heuristic Algorithm

The numerical experiments introduced later reveal that the column generation algorithm does not scale well to problems with a large number of entities. To overcome the scalability problem, we propose a heuristic method called the CG Heuristic algorithm that relies on column generation.

The CG Heuristic algorithm first finds a partition with a large number of clusters by neglecting the cardinality constraint. In the second step, we combine the clusters by considering unions to obtain exactly K clusters while obeying the cardinality constraint, which is a slight variant of column generation. We refer

Algorithm 2 $CG(\delta^{(0)}, \xi^{(0)}, k^{max})$

```
1:  $k \leftarrow 0$ , randomly generate  $\bar{\mathcal{S}}^{(0)}$  (a small subset of  $\mathcal{S}$ )
2: while  $k < k^{max}$  do
3:    $(z^{(k)}, q^-, q^+; \pi^{(k)}, v) \leftarrow$  solve the stabilized master problem
4:    $S^{(k)} \leftarrow$  solve the pricing problem with  $\pi^{(k)}$ 
5:   if  $a_S^T \pi^{(k)} + v \leq c_S$  and  $q^- = q^+ = \mathbf{0}$  then
6:      $z^* \leftarrow z^{(k)}$ , and stop
7:   else
8:      $\bar{\mathcal{S}}^{(k+1)} \leftarrow \bar{\mathcal{S}}^{(k)} \cup S^{(k)}$ 
9:     if  $a_S^T \pi^{(k)} + v \leq c_S$  then
10:       $\delta^{(k+1)} \leftarrow$  Update  $(\delta^{(k)})$ 
11:       $\xi^{(k+1)} \leftarrow$  Update  $(\xi^{(k)})$ 
12:    end if
13:  end if
14:   $k \leftarrow k + 1$ 
15: end while
```

to the intermediate clusters from the first part, which are the input to the column generation algorithm in the second part, as groups.

The algorithmic framework is presented in Algorithm 3. We require that $R > K$ since the second step is to combine R groups into K clusters. Lines 1-5 represent the first step to create R groups and Line 6 represents the second step to find a solution to the original problem. Lines 1-5 are basic and does not need further explanations. It yields R “low cost” groups. Since $R > K$, in Line 6 we combine some groups so that we end up with exactly K clusters, each one with cardinality at least n . This regrouping of groups is performed in an optimal way by using the column generation framework.

Algorithm 3 CG Heuristic (R)

Require: $R(> K)$

```
1: Randomly generate  $R$  groups
2: while there is an update in groups do
3:   Perform regression over each group  $r \in [R]$  to obtain regression coefficients  $\beta_r$ 
4:   For  $i \in [I]$ , reassign entity  $i$  to group  $r^*$ , where  $r^* = \arg \min_r \sum_{l=1}^L (y_{il} - \sum_{j=1}^J \beta_{rj} x_{ilj})^2$ 
5: end while
6: Execute CG by treating each group as entity
```

For Line 6, we need to revise the master and pricing problems in the following way when we cluster a group of entities instead of single entities. Suppose at the end of Line 5 we clustered I entities into R groups $\{G_1, G_2, \dots, G_R\}$, and then apply the column generation algorithm to the R groups of entities. Let \mathcal{S}_R be the set of all subsets S of $[R]$ such that $|\cup_{r \in S} G_r| \geq n$, and let $a_{rS} = 1$ if $r \in S$, and $a_{rS} = 0$ otherwise. To obtain the new master problem, we need to replace \mathcal{S} with \mathcal{S}_R and a_{iS} with a_{rS} in the master problem (6)–(8). In addition, the range of constraints (8) changes to $r \in [R]$.

We denote the dual variables of constraints (8) in the master problem by π_r , and introduce the binary decision variables z_r for $r \in [R]$ to indicate whether group G_r is selected in the cluster with the minimum reduced cost. To obtain the new pricing problem, we need to replace z_i 's with z_r 's in the pricing problem (10)–(13). Constraint (13) is changed to $\sum_{r=1}^R |G_r| z_r \geq n$, and the range in constraints (11) and (12) now becomes $r \in [R]$, $i \in G_r$, $l \in [L]$. The new pricing problem has the same number of constraints as the pricing problem (10)–(13), however, it has only R binary variables, comparing to I such variables in the pricing problem (10)–(13).

3.3 GA-Lloyd Heuristic Algorithm

Scientific works as those presented by Maulik and Bandyopadhyay [18] and Chang et al. [5], effectively suggest to embed the concept of the Lloyd’s algorithm, which is sometimes referred to as the k-means algorithm, into a genetic search metaheuristic framework to find proper clusters for the MSSC clustering problem. Here we discuss our proposed adaptation of the GA-based Lloyd’s clustering algorithms for solving the generalized CLR problem.

For the Lloyd’s algorithm part, a vector of regression coefficients β_k is used to represent cluster k , and an entity is recursively assigned to the cluster that gives the smallest sum of squared errors for this entity. The GA part helps escape local optimal solutions. The overall algorithmic framework is outlined in Algorithm 4.

Algorithm 4 GA-Lloyd

Require: K, \maxIter, H, p_m

- 1: For each h in $[H]$, create $C(h)$ by randomly generating K clusters
 - 2: **while** objective function value improved in the previous \maxIter iterations **do**
 - 3: Randomly select parent chromosomes h_1 and h_2 using roulette wheel selection
 - 4: Create child chromosomes h_a and h_b by performing crossover on h_1 and h_2
 - 5: Mutation on $h \in \{h_a, h_b\}$
 - 6: Obtain $C(h_a)$ and $C(h_b)$ based on Lloyd’s algorithm, calculate $\gamma(h_a)$ and $\gamma(h_b)$
 - 7: **if** $\max\{\gamma(h_a), \gamma(h_b)\} < \min_{h \in [H]} \gamma(h)$ **then**
 - 8: Replace h^* , $h^* = \operatorname{argmin}_{h \in [H]} \gamma(h)$, with the chromosome with larger fitness among h_a and h_b
 - 9: **end if**
 - 10: **end while**
-

In Line 1, we start by randomly generating H partitions $C(1), \dots, C(H)$, each of which corresponds to K clusters of entities with $C(h) = \{C_1(h), C_2(h), \dots, C_K(h)\}$ for $h \in [H]$. Any randomly generated partition $C(h)$ has to satisfy the constraint that $|C_k(h)| \geq n$ for $k \in [K]$. A population P consists of H chromosomes, and chromosome h is encoded as a vector $\beta(h)$ of size $J \cdot K$. In any chromosome, the first J genes represent the regression coefficients $\beta_1(h)$ for the first cluster, and the next J genes represent the regression coefficients $\beta_2(h)$ for the second cluster, and so on. The encoding of chromosome h is illustrated in Figure 1.




Figure 1: Encoding of Chromosome h

The regression coefficient $\beta_k(h)$ is obtained by running regression over cluster $C_k(h)$. The fitness $\gamma(h)$ of the chromosome h is defined to be

$$\gamma(h) = \frac{1}{\sum_{k=1}^K \sum_{i \in C_k(h)} \sum_{l=1}^L (y_{il} - \sum_{j=1}^J \beta_{kj}(h) x_{ilj})^2}. \quad (14)$$

We continue by performing the following genetic operations on the population of chromosomes iteratively until the number of iterations without improvement reaches a specified maximum number \maxIter . First, in Line 3, we randomly select two parent chromosomes h_1 and h_2 from population P using roulette wheel selection. Chromosome h is chosen with probability $\gamma(h) / \sum_{g=1}^H \gamma(g)$. Second, in Line 4, we perform crossover on chromosomes h_1 and h_2 . We select a gene position as a random integer in the range of $[1, K \cdot J - 1]$. We require this random integer to be no more than $K \cdot J - 1$ so that there is at least one gene positioned to the right of it. The portions of the chromosome lying to the right of this gene position are exchanged to produce two child chromosomes h_a and h_b encoded by $\beta(h_a)$ and $\beta(h_b)$. The crossover operation is illustrated in Figure 2.

Third, in Line 5, we perform mutation on these two child chromosomes. The mutation is performed on a child chromosome with a fixed probability p , where p is a parameter. A gene position with value v is




Figure 2: Crossover of Parent Chromosomes h_1 and h_2

randomly picked from the child chromosome. After mutation, it is changed to $v \pm 2\delta v$ with equal probability if v is not zero. Here δ is a random number with uniform distribution between zero and one. Otherwise, when v is zero, it is changed to $v \pm 2\delta$ with equal probability. In this way, the regression coefficients can take any real values after sufficient number of iterations. Next, in Line 6, we need to decode these two mutated child chromosomes to get the partitions $\mathcal{C}(h_a)$ and $\mathcal{C}(h_b)$ of clusters they represent. To decode the child chromosome h_a , we assign entity i to cluster $C_{k_i^*}(h_a)$ for

$$k_i^* = \arg \min_k \sum_{l=1}^L (y_{il} - \sum_{j=1}^J \beta_{k_j}(h_a) x_{ilj})^2$$

Then, we perform regression over each cluster of $\mathcal{C}(h_a)$ and $\mathcal{C}(h_b)$, and update the encoding of these two child chromosomes $\beta(h_a)$ and $\beta(h_b)$ with the resultant regression coefficients. Fitness $\gamma(H+1)$ and $\gamma(H+2)$ are calculated for the child chromosomes using (14). In Lines 6-7, we replace the chromosome in population P with the smallest fitness with the child chromosome with the smaller fitness if $\max(\gamma(h_a), \gamma(h_b)) < \min_{h \in [H]} \gamma(h)$.

During the decoding step of the GA-Lloyd algorithm, we may need to adjust the clusters generated in order to satisfy the minimum size constraints. If cluster C_i has size smaller than n , then we sort the entities not in C_i in the increasing order of the sum of squared regression errors, and then reassign these entities in the sorted order to cluster C_i until the size of C_i reaches n . We also skip each entity that would reduce the size of its original cluster below n . When there is more than one cluster with size smaller than n , we perform this adjustment for the smallest cluster first.

3.4 Two-Stage Heuristic Algorithm for SKU Clustering Problem

Due to its simplicity, two-stage heuristic algorithms are frequently employed in practice for solving the CLR problem. In the first stage, entities are partitioned according to certain approximate measures of the regression coefficients. In the second stage, regression models are built over the resultant clusters. The clustering method for the first stage is usually problem specific.

In this section, we describe our two stage heuristic algorithm for the SKU clustering problem. The first stage of our algorithm is based on hierarchical clustering and simple (one dimensional) regression. In detail,

- (1) We carry out one regression for each SKU to get sales without promotional effects. The dependent variable for the regression is the weekly sales, and the independent variable is the price discount. Note that these regressions are one dimensional.

- (2) We construct, for each SKU, a sales vector of dimension 52 (i.e. the number of weeks in a year). The l^{th} element of the vector records the mean sales without promotional effects of the l^{th} week of the year, averaged over all years. The mean sales without promotional effects are equal to the actual sales minus the promotional effects obtained in the previous step.
- (3) We calculate the correlations between any pair of sales vectors constructed in the previous step. The distance between any two SKUs is defined as one minus their corresponding correlation.
- (4) We perform agglomerative hierarchical clustering over SKUs using distances generated in the previous step. We use the maximum distance between SKUs of each cluster as the distance between any two clusters as in the complete linkage clustering [15].

4 Numerical Experiments

In this section, we examine the performance of CG (Algorithm 2), CG Heuristic (Algorithm 3), GA-Lloyd (Algorithm 4) and the two-stage algorithms on the SKU clustering problem according to its seasonal effects. The regression model for this problem has the following form:

$$f_0(\text{weekly sales}) = f_1(\text{promotional predictors}) + f_2(\text{seasonal predictors}),$$

where the seasonal effects are modeled by 52 dummy variables, one for each week. There are two types of data sets we used in the experiment.

1. Real-world data: We use real-world data from a large retail chain. We omit the exact form of the regression model due to confidentiality. It includes more than two years of aggregated sales and promotional data of the entire chain. The products within this chain are grouped into subcategories for purchasing purposes. However, it is assumed that the products within the same subcategory have different seasonal patterns with regard to promotions. We tested our algorithms on two representative subcategories from the data, a smaller subcategory “Cream” and one of the largest subcategories “Medicines.” Both subcategories have more than one seasonal pattern. Each SKU in these subcategories has at least 52 weeks of data and at most 129 weeks of data from year 2006 to 2008.
2. Synthetic data: We generate random instances that have similar patterns as the real-world data. The promotional predictor includes percentage of discount and the seasonal predictor captures week index. Therefore, the data set includes weekly sales, percentage of discount, and the week index. Each entity has a year of records (52 weeks). For each $I \in \{15, 20, 25, 50, 100, 150, 200\}$, we generate 10 random instances, which results into 7 different size data sets of total 70 instances. The detailed instance generation procedure is available in Appendix B and the data set is available at a web site ¹.

The following computational environments were used.

1. Real-world data: All the algorithms except the two-stage algorithm were implemented in Java 1.6 with CPLEX 11 as the mathematical programming solver. The “lm” function in R 2.8 is used to perform regressions for the GA-Lloyd algorithm. The two-stage algorithm is implemented in R 2.8, and the “hclust” function for hierarchical clustering is employed to perform clustering. All numerical experiments were performed on a 64-bit server with a multi-core Intel Xeon 2 GHz CPU and 10 GB of RAM.
2. Synthetic data: All the algorithms were implemented in Java 1.7 with CPLEX 12.5 and R 2.8. The experiments were performed on a 64-bit server with a multi-core Intel Xeon 2.8 GHz CPU and 32 GB of RAM.

Note also that, for all experiments in Section 4, specific values for the following parameters of the GA-Lloyd algorithm were selected based on a sensitivity analysis: $H = 10$ (population size) and $p = 0.01$ (permutation probability). Furthermore, as a termination criterion for the algorithm we specify that the number of consecutive iterations with no improvement has to reach 50.

¹dynresmanagement.com/public_data/sku_clustering_random_data.zip

4.1 Comparison of the Heuristic Algorithms

4.1.1 Real-World Data

Based on preliminary computational studies of the column generation algorithm and the results shown in Section 4.4, we observed that the algorithm does not scale well when applied to large size instances. Hence, the heuristic algorithms are employed to cluster SKUs for the “Cream” and “Medicine” subcategories.

In this section we compare the performance of the GA-Lloyd, two-stage, and CG Heuristic algorithms with $K \in \{2, 3, 4, 5\}$ and $n = 3$ in terms of solution time and quality. As the initial important remark, we observe that the GA-Lloyd algorithm performs the best while providing a good balance between time and quality. In the experiments, an instance denoted by “66_2” implies that we divided 66 SKUs into two clusters. The stopping criteria for the CG Heuristic algorithm are: (a) an optimal solution is found by the column generation with R groups of entities, and (b) a time limit of 10 hours is reached after the last pricing problem. For subcategory “Cream,” the CG Heuristic algorithm applied stopping criterion (a), whereas for the subcategory “Medicine,” it stopped due to (b).

First, we compare the GA-Lloyd algorithm with the two-stage algorithm. Figure 3 shows the relative improvement (RI%) of the GA-Lloyd algorithm over the two-stage algorithm, which is given by

$$RI = \frac{SSE(\text{Two Stage}) - SSE(\text{GA-Lloyd})}{SSE(\text{GA-Lloyd})},$$

where the value of $SSE(\text{algorithm})$ is the sum of the regression errors across all clusters. We observe that a substantial improvement is achieved by the GA-Lloyd algorithm over the two-stage heuristic.

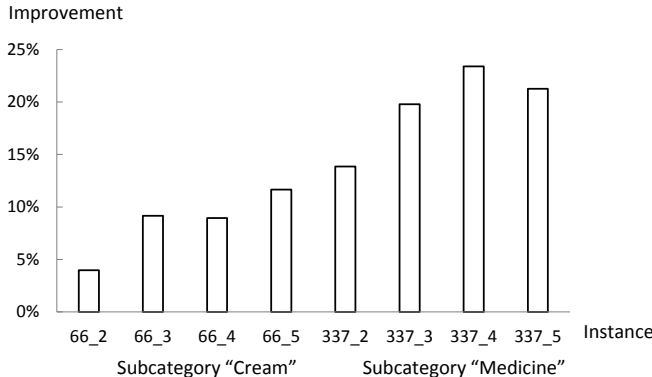


Figure 3: Percentage Improvement of GA-Lloyd over Two-Stage for Real-World Data

In terms of the running time, we observe that the two-stage algorithm outperforms the GA-Lloyd algorithm while converging within one and five minutes, respectively, for the “Cream” and “Medicine” subcategories, whereas the GA-Lloyd algorithm took roughly 20 minutes and one hour for the corresponding subcategories. In practice, however, run times of one-hour are acceptable for the SKU clustering problem according to our retail partner.

We also compare the GA-Lloyd algorithm with the CG Heuristic algorithm on the two subcategories. We define the percentage cost difference as

$$\text{CostDiff} = \left[\frac{SSE(\text{GA-Lloyd}) - SSE(\text{CG Heuristic})}{SSE(\text{GA-Lloyd})} \right],$$

where $\text{CostDiff} > 0$ means that the CG Heuristic algorithm generates a better solution and $\text{CostDiff} < 0$ implies the superiority of the GA-Lloyd solution.

Figure 4 shows the percentage cost difference between the GA-Lloyd algorithm and the CG Heuristic algorithm with $R = \{6, 8, 10\}$. We observe that the resulting CostDiff values between the GA-Lloyd and CG Heuristic algorithms are not significant, all within 6%. The CG Heuristic algorithm with $R = 10$ generates better solutions than the GA-Lloyd algorithm for five out of eight instances. Among the five instances,

in four of them the improvement is barely noticeable (“337_2” is the only case with more pronounced improvement). However, its running time, four to six hours for subcategory “Cream” and 10 hours for subcategory “Medicine”, is much longer than that of the GA-Lloyd algorithm, as illustrated in Figure 5. In summary, it is recommended to select largest R within affordable time limit for solving the problem.

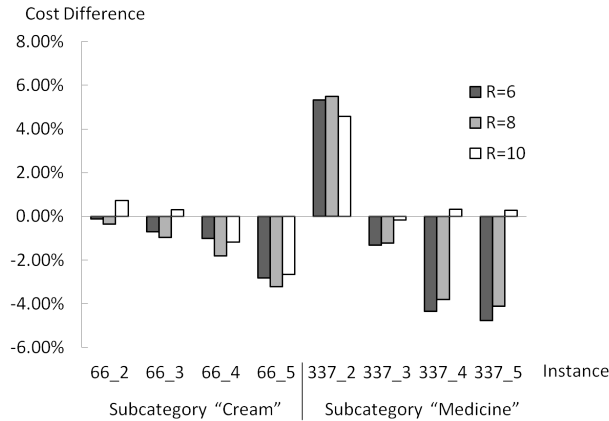


Figure 4: Percentage Cost Difference between GA-Lloyd and CG Heuristic for Real-World Data

From Figure 4, we also observe that the GA-Lloyd algorithm performs better than the CG Heuristic algorithm with $R = \{6, 8\}$ for all but one instance. In addition, the GA-Lloyd algorithm also outperforms the CG Heuristic algorithm in terms of computational times for all these instances, as shown in Figure 5. When comparing the solutions of the CG Heuristic algorithm with $R = 6$ and 8 for subcategory “Cream”, Figure 4 shows that a higher value of R does not necessarily imply a better solution.




Figure 5: CG Heuristic Running Time for Real-World Data

4.1.2 Synthetic Data

In this section, we compare the performance of the CG Heuristic and GA-Lloyd algorithms with $K \in \{2, 3, 4\}$ and $n = 3$ for the synthetic instances with $I \in \{50, 100, 150, 200\}$. Due to its excessive computational time, we did not run the CG algorithm. We execute the CG Heuristic algorithm with $R = 8$. The result is presented in Figure 6. Figure 6(a) shows the percentage cost difference between the GA-Lloyd and CG Heuristic algorithms. Similar to the result for the real-world data, we observe that the CG Heuristic algorithm performs better when K is small, and the GA-Lloyd algorithm starts to outperform as K increases. In terms of the running time, we observe that the two algorithms are of the same magnitude as presented in Figure 6(b).




Figure 6: Comparison of CG Heuristic and GA-Lloyd for Synthetic Data

4.2 Seasonal Patterns Identified by GA-Lloyd

Based on the superiority of the GA-Lloyd algorithm over its counterparts shown in the previous section, in this section we present numerical results of the seasonal patterns identified by this algorithm when applied to the largest subcategory, “Medicine”.

Figure 7 shows the seasonal multipliers obtained from the GA-Lloyd algorithm for each cluster when dividing the subcategory “Medicine” into 2, 3, 4 and 5 clusters. From the figure, we observe distinct seasonal patterns: (1) U-shaped curve, (2) inverted V-shaped, and (3) flat. Observe that all of the three seasonal patterns have been found when dividing SKUs into three clusters. This observation indicates that it is not necessary to further divide them into four or five clusters since some seasonal patterns look similar. The three clusters of SKUs represent medicines that intuitively have such different seasonal patterns: one corresponding to medicines (such as for cold and flu) that sell more in the winter, one corresponding to medicines (such as for bug repellents and sunburns) with uplift in the summer, and one corresponding to medicines (such as for diarrhea and constipation) with stable sales year around.

4.3 Optimality Gap of GA-Lloyd and CG Heuristic

In this section, we benchmark the performance of the GA-Lloyd and CG Heuristic algorithms by comparing against the CG algorithm, which is an optimal algorithm. In order to measure the performance, we calculate

$$OptGap = \left[\frac{SSE(algo) - SSE(Column\ Generation)}{SSE(Column\ Generation)} \right],$$

where $algo \in \{ CG\ Heuristic, GA-Lloyd \}$. For the real-world data, note that since the column generation algorithm did not cluster SKUs for the “Cream” and “Medicine” subcategories within a reasonable computational time due to their large sizes, we construct smaller instances with 15 and 20 SKUs that are randomly chosen from the large subcategory “Medicines.” For these instances, we test the algorithms for parameters $K \in \{2, 3, 4\}$ and $n = 2$. We also studied the instance that divides the subcategory “Cream” with 66 SKUs into two clusters with minimum cluster size $n = 3$, which we refer to as 66₂ instance in this section. The column generation exact algorithm were executed with 10 hours of time limit for the 66₂ instance. For the synthetic data, we test small instances with $I \in \{15, 20, 25\}$ for parameters $K \in \{2, 3, 4\}$ and $n = 2$.

4.3.1 Real-World Data

Figure 8(a) shows the optimality gap values of the GA-Lloyd algorithm. We observe that the GA-Lloyd algorithm achieves close to optimal solutions, with optimality gaps less than 2%. In addition, the GA-Lloyd algorithm finishes within five minutes for these smaller instances.

For the 66₂ instance, the solution obtained from the column generation exact algorithm after 10 hours of running time is only 1.47% better than the GA-Lloyd solution, which is obtained in less than 20 minutes. This case is not shown in the figure.

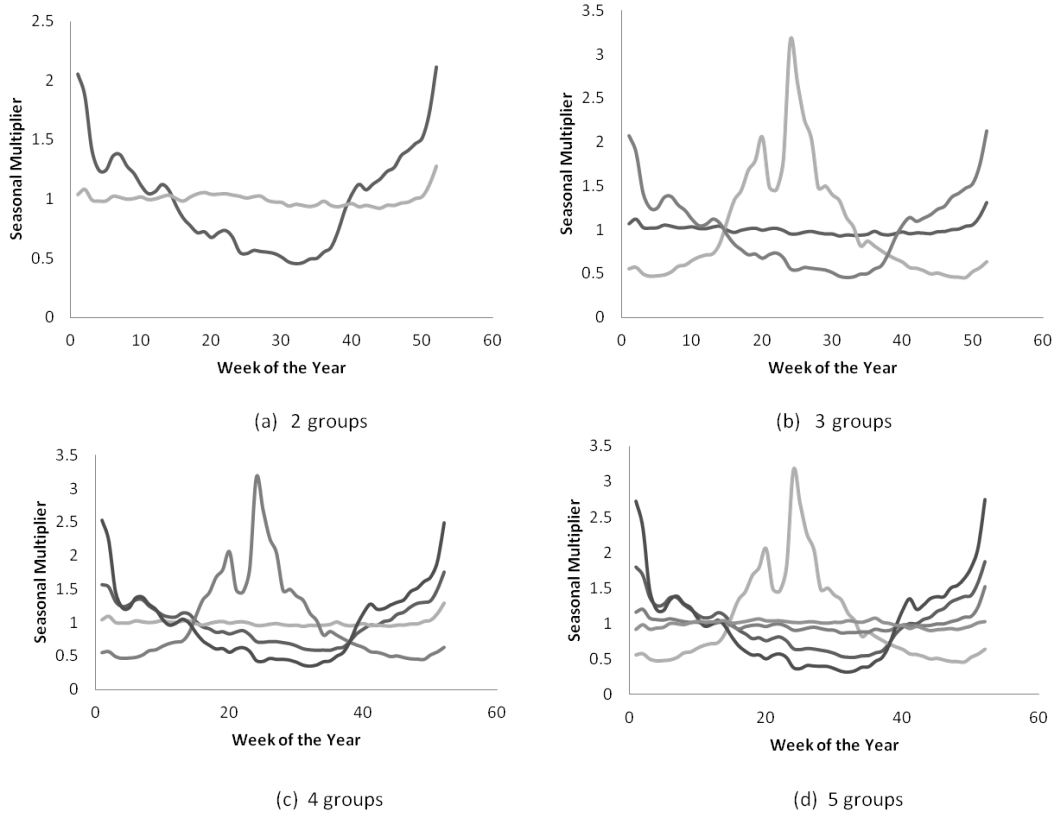


Figure 7: Subcategory “Medicine” Seasonal Multipliers for Real-World Data

Figure 8(b) shows the optimality gap values of the CG Heuristic algorithm for different values of R . We observe that the the CG Heuristic algorithm also finds close to optimal solutions, with optimality gaps less than 5%. In addition, the CG Heuristic algorithm finishes within five minutes for the instances of size 15, and within 20 minutes for the instances of size 20. By comparing the solutions for the instance “20_2” for R equal 8 and 10, we again find that a larger R does not necessary generate a better solution.

For the 66₂ instance, the solution obtained from column generation exact algorithm after 10 hours of running time is only 0.13%, 1.86%, 0.75% better than the CG Heuristic solution with R equal to 6, 8, and 10, respectively. The corresponding running times of the CG Heuristic algorithm are 54, 178, and 365 minutes, respectively. The performance for the 66 SKUs is not depicted graphically.

The average gap in Figure 8(a) is 1.14%, while in Figure 8(b) for $R = 10$ it is 1.01%. This indicates that for smaller instances the CG Heuristic algorithm outperforms the GA-Lloyd algorithm if the objective value is the only performance indicator and the computational time is limited to 10 hours. On the contrary, Figure 4 indicates clearly that the GA-Lloyd algorithm suits better for larger instances. This implies that whenever a strict run time limit is imposed, the GA-Lloyd algorithm is very likely to outperform its counterparts for most of the instances.

4.3.2 Synthetic Data

In this section, we test the performance of the CG Heuristic and GA-Lloyd algorithms by comparing against the solution of CG for the small synthetic instances with $I \in \{15, 20, 25\}$ and $K \in \{2, 3, 4\}$. We only execute the CG Heuristic algorithm with $R = 8$. In Figure 9(a), we observe that the running time of the algorithms are of the same magnitude. We also observe that the running time of the CG Heuristic algorithm tends to decrease in K and the running time of the GA-Lloyd algorithm tends to increase in K . For both algorithms, the optimality gaps are generally larger than the result from real-world data in Section 4.1.1. The optimality

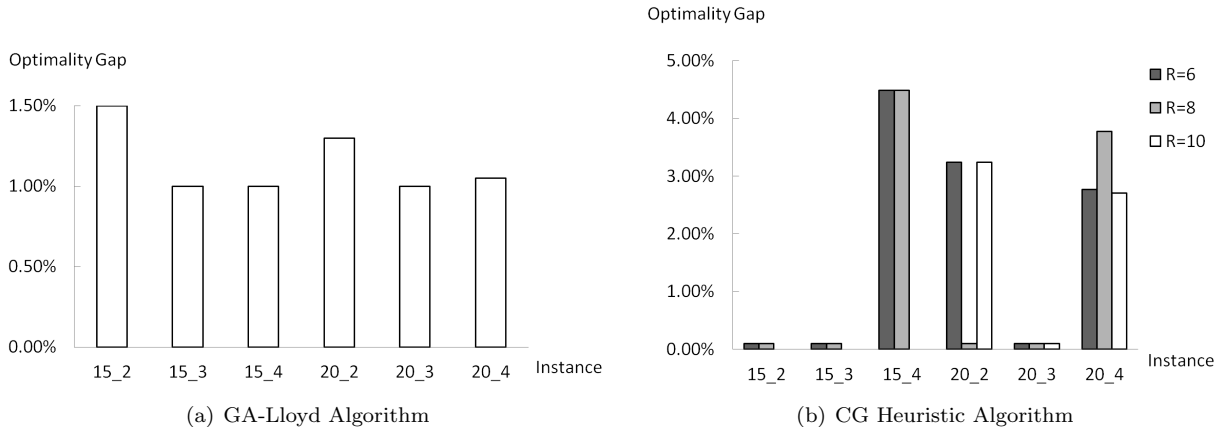


Figure 8: Optimality Gap of GA-Lloyd and CG Heuristic for Real-World Data

gap in Figure 9(b) shows that the GA-Lloyd algorithm performs better than the CG Heuristic algorithm.




Figure 9: Comparison of CG Heuristic and GA-Lloyd for Synthetic Data

4.4 Time Study of the Column Generation Algorithm

4.4.1 Real-World Data

The performance of the Column Generation algorithm (Algorithm 2) is assessed on a set of computational experiments conducted on instances with 15, 20 and 66 SKUs. These instances are chosen in a similar way as in the previous section. Figure 10(a) presents the running time of Algorithm 2 for specific numbers of SKUs and clusters to divide in. We can observe that it takes roughly 3 hours to divide 20 SKUs into two clusters.

For the instance with 66 SKUs and two clusters, we are unable to get an optimal solution after 10 hours. When comparing to the lower bound obtained by solving the linear relaxation of the mixed integer formulation (1)–(5), the solution gotten from column generation after 10 hours of running time is 38.23% larger than the lower bound. However, we suspect this solution is very close to the optimal one because we observe that the minimum reduced cost of the master problem is close to zero.

We also study another version of the generalized CLR problem with the sum of absolute errors as the objective to examine whether the difficulty in solving the pricing problem is due to the nonlinearity of the objective function in the pricing problem (10)–(13). More specifically, we change the objective function in the

pricing problem to $\sum_{i=1}^I \sum_{l=1}^L t_{il} - \sum_{i=1}^I \pi_i z_i$. Figure 10(b) presents the running time when the objective function for the CLR problem is the sum of absolute errors. From this figure, we observe that the running time also increases dramatically as the number of SKUs increases. It also takes hours to solve the instances with 20 SKUs. Therefore, we believe the nonlinear objective of the pricing problem is not the complicating factor that greatly drives up the computation time. These running times are higher due to a larger number of iterations resulting from degeneracy of LP solutions (the per iteration time is lower).

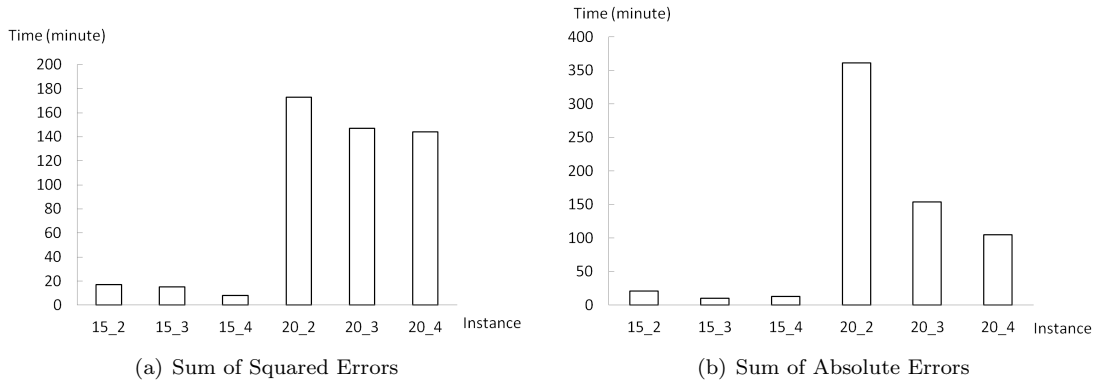


Figure 10: Column Generation Running Time for Real-World Data

4.4.2 Synthetic Data

We present the running time of the CG algorithm for the synthetic data. The running times are of a different magnitude from the one for the real-world data. This is because (1) the computational environments are different (2) and the data has smaller L and attributes J . In Figure 11, we plot the running time of CG. We observe that the running time drastically increases as I increases. The running time decreases in K with fixed I .




Figure 11: Column Generation Running Time for Synthetic Data

5 Conclusions

We propose an exact column generation algorithm, the CG Heuristic algorithm, the GA-Lloyd metaheuristic, and the two-stage algorithm for the resolution of the generalized cluster-wise linear regression problem. We examine the performance of our algorithms on the SKU clustering problem according to seasonal effects using a real-world retail data set from a large retail chain. We find that the column generation exact algorithm can solve small instances to optimality. We use the column generation exact algorithm to benchmark the

performance of the GA-Lloyd algorithm and the CG Heuristic algorithm, although the CG algorithm cannot scale well to instances of large sizes. The two-stage algorithm can produce SKU clusters very fast, but with higher objective values than the GA-Lloyd algorithm. The CG Heuristic algorithm performs slightly better than the GA-Lloyd algorithm for some instances, but with much longer running time. The GA-Lloyd algorithm provides a good balance between solution quality and time, and generates SKU clusters with distinctive seasonal patterns efficiently and effectively.

Acknowledgement: We are much obliged to Mr. Molham Aref, the CEO of Predictix Inc, for his assistance during the project. Mr. Aref allowed us to use the data and he approved a summer internship by Yan Jiang.

References

- [1] Aloise, D., Deshpande, A., Hansen, P., and Popat, P. (2009). NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning*, 75:245–248.
- [2] Aloise, D., Hansan, P., and Liberti, L. (2012). An improved column generation algorithm for minimum sum-of-squares clustering. *Mathematical Programming Series A*, 131:195–220.
- [3] Barnhart, C., Johnson E. L., Nemhauser G. L., Savelsbergh, M.W.P., and Vance, P.H. (1998). Branch-and-Price: Column Generation for Solving Huge Integer Programs. *Operations Research*, 46:316–329.
- [4] Bertsimas, D. and Shioda, R. (2007). Classification and regression via integer optimization. *Operations Research*, 55:252–271.
- [5] Chang, D., Zhang, X., and Zheng, C. (2009). A genetic algorithm with gene rearrangement for k-means clustering. *Pattern Recognition*, 42:1210–1222.
- [6] DeSarbo, W. S. (1989). A simulated annealing methodology for clusterwise linear regression. *Psychometrika*, 54:707–736.
- [7] DeSarbo, W. S. and Cron, W. L. (1988). Regression clustering. *Journal of Classification*, 5:249–282.
- [8] du Merle, O., Hansen, P., Jaumard, B., and Mladenović, N. (2000). An interior point algorithm for minimum sum-of-squares clustering. *SIAM Journal on Scientific Computing*, 21:1485–1505.
- [9] du Merle, O., Villeneuve, D., Desrosiers, J., and Hansen, P. (1999). Stabilized column generation. *Discrete Mathematics*, 194:229–237.
- [10] D’Urso, P., Massari, R., and Santoro, A. (2010). A class of fuzzy clusterwise regression models. *Information Sciences*, 180:4737–4762.
- [11] Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Company.
- [12] Hansen, P. and Jaumard, B. (1997). Cluster analysis and mathematical programming. *Mathematical Programming*, 79:191–215.
- [13] C. Hennig (1999). Models And Methods For Clusterwise Linear Regression. *Proceedings in Computational Statistics*.
- [14] Ingrassia, S., Minotti S.C., and Punzoa, A. (2014). Model-based clustering via linear cluster-weighted models. *Computational Statistics and Data Analysis*, 71:159–182.
- [15] Johnson, R. and Wichern, D. (2007). *Applied multivariate statistical analysis* . Pearson.
- [16] Lau, K., Leung, P., and Tse, K. (1999). A mathematical programming approach to clusterwise regression model and its extensions. *European Journal of Operational Research*, 116:640–652.

- [17] Luo, Z. and Chou, E. Y. (2006). Pavement condition prediction using clusterwise regression. In *Transportation Research Board 2006 Annual Meeting*.
- [18] Maulik, U. and Bandyopadhyay, S. (2000). Genetic algorithm-based clustering technique. *Pattern Recognition*, 33:1455–1465.
- [19] Muruzábal, J., Vidaurre, D., and Sánchez, J. (2012). SOMwise regression: a new clusterwise regression method. *Neural Computing and Applications*, 21:1229–1241.
- [20] Openshaw, S. (1977). A geographical solution to scale and aggregation problems in region-building, partitioning and spatial modeling. *Transactions of the Institute of British Geographers, New Series*, 2:459–472.
- [21] Spath, H. (1979). Clusterwise linear regression. *Computing*, 22:367–373.
- [22] Vanderbeck, F. and Wolsey, L. A. (1996). An exact algorithm for IP column generation. *Operations Research Letters*, 19:151–159.
- [23] Wedel, M. and DeSarbo, W. S. (1994). *Advanced methods of marketing research*, chapter A review of recent developments in latent structure regression models, pages 352–388. Blackwell Publishing, London.
- [24] Zhang, B. (2003). Regression clustering. In *Proceedings of the Third IEEE International Conference on Data Mining (ICDM03)*, pages 451–458.

APPENDIX

A Proof of Theorems

A.1 Proof of Theorem 1

Intuitively, the generalized CLR problem resembles the MSSC problem, which is known to be NP-hard. We conduct a polynomial transformation from MSSC to a special case of the CLR problem as follows.

Consider an instance of the MSSC problem with I entities. Each entity i has an associated vector $\mathbf{y}_i = (y_{i1}, y_{i2}, \dots, y_{iL})$. Let vector \mathbf{y}_i be the observations of the dependent variable for entity i , and let an identity matrix of size L be the observations of independent variables \mathbf{x}_i , which means

$$\mathbf{x}_i = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}. \quad (15)$$

This yields an instance of the generalized CLR problem of dimension $L \times L$. The regression coefficient β_k for cluster k is the centroid of the entities assigned to cluster k with

$$\beta_{kl} = \frac{\sum_{i \in C_k} y_{il}}{|C_k|}.$$

This proves NP-hardness of the generalized CLR problem. □

A.2 Proof of Theorem 2

In this section, we show that the pricing problem is NP-complete.

Let us consider a special case of the pricing problem with the observations of the independent variables \mathbf{x}_i being an $L \times L$ identity matrix as in (15). Then the pricing problem becomes

$$\min_{|S| \geq n, \beta} \sum_{i \in S} \sum_{l=1}^L (y_{il} - \beta_l)^2 - \sum_{i \in S} \pi_i.$$

Given any cluster S such that $|S| \geq n$, by equating the first order derivative of the pricing objective function to zero, we obtain the optimal $\beta(S)$ as the centroid of vector \mathbf{y}_i :

$$\beta_l(S) = \frac{\sum_{i \in S} y_{il}}{|S|}.$$

The Huygen's theorem states that for a given set S of vectors $\mathbf{u}_i = (u_{i1}, u_{i2}, \dots, u_{iL})$, the sum of squared distances to the centroid is equal to the sum of squared distances between these vectors divided by two times the cardinality of the set, which mathematically stated reads

$$\sum_{i \in S} \sum_{j \in S, j \neq i} \|\mathbf{u}_i - \mathbf{u}_j\|_2^2 = 2|S| \sum_{i \in S} \|\mathbf{u}_i - \bar{\mathbf{u}}(S)\|_2^2$$

where $\bar{u}_l(S) = \frac{\sum_{i \in S} u_{il}}{|S|}$ and $\bar{\mathbf{u}}(S) = (\bar{u}_1(S), \bar{u}_2(S), \dots, \bar{u}_L(S))$. Based on Huygen's theorem, this special case of the pricing problem can also be stated as:

$$\min_{|S| \geq n} \sum_{i \in S} \sum_{j \in S, j \neq i} \|\mathbf{y}_i - \mathbf{y}_j\|_2^2 - 2|S| \sum_{i \in S} \pi_i.$$

By using a transformation from the independent set problem, [11], we show that this special case of the pricing problem with this formulation is NP-complete, which implies the pricing problem is NP-complete.

Let us now formally prove the theorem. We first introduce the independent set problem ([11]), a known NP-complete problem, which is used to prove that the pricing problem is NP-complete.

Instance: Graph $G = (V, E)$, and a positive integer $M \leq |V|$;

Question: Does G contain an independent set of size M , i.e., a subset $V' \subseteq V$ with $|V'| = M$ such that no two vertices in $|V'|$ are joined by an edge in E ?

We next show, through a polynomial reduction from the independent set of size M problem that a constrained version of the pricing problem, which we refer to as "the subset of size M problem," is NP-complete. The subset of size M problem is as follows.

Instance: n vectors $(\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^n)$ of dimension m (i.e., $\mathbf{u}^i = (u_1^i, u_2^i, \dots, u_m^i)$), n real numbers π_i for $i \in [n]$, and another real number K ;

Question: Is there a subset $S \subseteq \{1, \dots, n\}$ of vectors with cardinality $|S| = M$ such that

$$\sum_{i \in S} \sum_{j \in S, j \neq i} \|\mathbf{u}^i - \mathbf{u}^j\|_2^2 - \sum_{i \in S} \pi_i \leq K?$$

Lemma 3 *The subset of size M problem is NP-complete.*

Proof: We show NP-completeness of the subset of size M problem using its relationship with the independent set of size M problem. Consider an instance of the independent set of size M problem with graph $G = (V, E)$. To each node $i \in V$, we assign a vector \mathbf{u}^i of size $|E|$. For $j = 1, \dots, |E|$, we have

$$u_j^i = \begin{cases} 1, & \text{if edge } (i, j) \in E \text{ and } i < j; \\ -1, & \text{if edge } (i, j) \in E \text{ and } i \geq j; \\ 0, & \text{otherwise.} \end{cases}$$

Let k_i be the degree of node i . If node i is connected to node j by edge $(i, j) \in E$, then

$$\|\mathbf{u}^i - \mathbf{u}^j\|_2^2 = (k_i - 1) + (k_j - 1) + (1 - (-1))^2$$

$$= k_i + k_j + 2,$$

and otherwise,

$$\|\mathbf{u}^i - \mathbf{u}^j\|_2^2 = k_i + k_j.$$

Let $\pi_i = 2(M - 1)k_i$ and $K = 0$. We next show that to answer the question whether there is a subset $V' \subseteq V$ with $|V'| = M$ such that $\sum_{i \in V'} \sum_{j \in V', j \neq i} \|\mathbf{u}^i - \mathbf{u}^j\|_2^2 - \sum_{i \in V'} \pi_i \leq 0$ is equivalent to answering the question whether there is an independent subset of size M .

If there is an independent subset $V' \subseteq V$ with $|V'| = M$, then $\sum_{i \in V'} \sum_{j \in V', j \neq i} \|\mathbf{u}^i - \mathbf{u}^j\|_2^2 - \sum_{i \in V'} \pi_i = \sum_{i \in V'} 2(M - 1)k_i - \sum_{i \in V'} 2(M - 1)k_i \leq 0$. If there does not exist an independent subset $V' \subseteq V$ with $|V'| = M$, then $\sum_{i \in V'} \sum_{j \in V', j \neq i} \|\mathbf{u}^i - \mathbf{u}^j\|_2^2 - \sum_{i \in V'} \pi_i \geq \sum_{i \in V'} 2(M - 1)k_i + 2 \cdot 2 - \sum_{i \in V'} 2(M - 1)k_i > 0$. Here the first inequality is because there are at least two nodes that are connected by an edge belonging to the subset V' . \square

We now show NP-completeness of the pricing problem by polynomially transforming the subset of size M problem to this problem. The decision version of our pricing problem is as follows.

Instance: I vectors $(\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^I)$ of dimension L (i.e., $\mathbf{u}^i = (u_1^i, u_2^i, \dots, u_L^i)$), I real numbers π_i for $i \in [I]$, a positive integer $n \leq I$, and another real number K .

Question: Is there a subset $S \subset \{1, \dots, I\}$ such that $\sum_{i \in S} \sum_{j \in S, j \neq i} \|\mathbf{u}_i - \mathbf{u}_j\|_2^2 - \sum_{i \in S} 2|S|\pi_i \leq K$?

If we can answer the pricing problem under the additional constraint that $|S| = M$ for $M = n, \dots, I$ in polynomial time, then we can answer the original pricing problem in polynomial time. The pricing problem with the additional constraint $|S| = M$ is the subset of size M problem with the same vectors \mathbf{u}^i and I real numbers $2|S|\pi_i$. Since the subset of size M problem is NP-complete, so is the pricing problem. \square


B Random Instance Generation

We generate random instances for the SKU clustering problem based on the following regression model.


$$\begin{aligned} \text{weekly sales} &= \text{regular and promotional sales} + \text{seasonal sales} + \text{random noise} \\ &= \text{discount} * \beta_{\text{discount}} + t \cdot \beta_t + \varepsilon, \end{aligned}$$

where t is the week index, β_t is the regression coefficient for week t , and β_{discount} is the regression coefficient for discount. In order to replicate realistic seasonal effects, we use seven equations for the coefficients of the seasonal effect, motivated from the result in Section 4.2. In Figure 12(a), the seven equations for the seasonal effect are presented. The horizontal axis represents the week number and the vertical axis represents the ratio between regular / promotional sales and average sales.

In this section, we denote a uniform random number with lower bound lb and upper bound ub as $U(lb, ub)$. We denote a normal random number with mean avg and standard deviation std as $N(avg, std)$.



(a) 7 patterns



(b) Example of generated instance with 15 entities

Figure 12: Seasonal effect patterns and simulated patterns

We present the overall instance generation algorithm in Algorithm 5. We generate 52 (weeks) time series data points for each identity i . The attributes of the final data set includes entity index (i), week number (t), weekly sales in t , and percentage of discount.

In Algorithm 5, we are given I and fixed $L = 52$. For each $i \in \{1, \dots, I\}$, we generate weekly sales and percentage of discount for 52 weeks based on the following steps. In Line 2, we first randomly generate average sales volume S_A using uniform distribution between 100 and 200. Then in Line 3, we randomly pick a seasonal function from the seven equations illustrated in Figure 12(a). Next in Line 4, $U(3, 6)$ weeks are randomly picked to be the weeks with promotion, where the remaining not selected weeks are without promotion. The discount attribute is generated for promotional weeks by randomly picking among 15%, 20%, 25%, and 30%. In Lines 5-11, for each $t \in [L]$, we generate sales by summing promotional and seasonal sales, and a random error specified below. In detail, in Lines 6-7, we generate regular and promotional sales. For weeks with promotion, the regular and promotional sales are $S_A \cdot (1 + p_{promo})$. If discount is 15%, then p_{promo} is $U(0.4, 0.5)$. If discount is 20%, 25%, and 30%, then p_{promo} is drawn from $U(0.5, 0.6)$, $U(0.6, 0.7)$, and $U(0.7, 0.8)$, respectively. For weeks without promotion, $p_{promo} = 0$ and regular and promotional sales are S_A . In Lines 8, seasonal sales D_s is obtained by multiplying $f_S(t)$ and S_A . In Line 9, random error ε is generated from normal distribution with zero mean and deviation $S_A/5$. Finally in Line 10, weekly sales in t are generated by summing D_p , D_s , and ε . In Figure 12(b), we plot example sales records for 15 entities.

Algorithm 5 Data generation

Require: I (number of entities or SKUs), $L = 52$

- 1: **for** $i = 1, \dots, I$
 - 2: Generate average demand $S_A \sim U(100, 200)$
 - 3: Randomly pick seasonal function $f_S(t)$
 - 4: Randomly pick $U(3, 6)$ promotional weeks and generate discount attribute
 - 5: **for** $t = 1, \dots, L$
 - 6: **if** t is promotional week, generate $D_p = S_A \cdot (1 + p_{promo})$
 - 7: **else** generate $D_p = S_A$
 - 8: Calculate $D_s = S_A \cdot f_S(t)$
 - 9: Generate $\varepsilon \sim N(0, \frac{S_A}{5})$
 - 10: $y_{it} = D_p + D_s + \varepsilon$
 - 11: **end for**
 - 12: **end for**
-

C Restricted Master Problem

In this section, we present a stabilized version of the master problem (6) - (8), referred to as restricted master problem, by applying the technique of du Merle et al. [9]. For iteration k , the restricted master problem is written as

$$\min \sum_{S \in \mathcal{S}} c_S z_S - \delta^{(k)} \mathbf{q}^- + \delta^{(k)} \mathbf{q}^+ \quad (16)$$

$$\sum_{S \in \mathcal{S}} z_S = K \quad (17)$$

$$\sum_{S \in \mathcal{S}} a_{iS} z_S - \mathbf{q}^- + \mathbf{q}^+ = 1 \quad i \in [I] \quad (18)$$

$$0 \leq \mathbf{q}^- \leq \boldsymbol{\xi}^{(k)} \quad (19)$$

$$0 \leq \mathbf{q}^+ \leq \boldsymbol{\xi}^{(k)} \quad (20)$$

$$z_S \in \{0, 1\} \quad S \in \mathcal{S},$$

which is obtained by introducing perturbation variables \mathbf{q}^- and \mathbf{q}^+ and stabilization parameters $\delta^{(k)}$ and $\boldsymbol{\xi}^{(k)}$.