

Predictive and Prescriptive Analytics for Location Selection of Add-on Retail Products

Teng Huang^{*1}, David Bergman^{†1}, and Ram.Gopal^{‡1}

¹Department of Operations and Information Management, School of Business, University of Connecticut

Abstract

In this paper, we study an analytical approach to selecting expansion locations for retailers selling add-on products whose demand is derived from the demand of another base product. Demand for the add-on product is realized only as a supplement to the demand of the base product. In our context, either of the two products could be subject to spatial autocorrelation where demand at a given location is impacted by demand at other locations. Using data from an industrial partner selling add-on products, we build predictive models for understanding the derived demand of the add-on product and establish an optimization framework for automating expansion decisions to maximize expected sales. Interestingly, spatial autocorrelation and the complexity of the predictive model impact the complexity and the structure of the prescriptive optimization model. Our results indicate that the models formulated are highly effective in predicting add-on product sales, and that using the optimization framework built on the predictive model can result in substantial increases in expected sales over baseline policies.

Keywords: Add-on products; Derived demand; Empirical demand estimation; Retail expansion optimization; Predictive and prescriptive analytics

*Teng.Huang@uconn.edu

†David.Bergman@uconn.edu

‡Ram.Gopal@uconn.edu

1 Introduction

Managers of retail chains face constant pressure to increase their sales. Given that the geographic reach of one retail store is necessarily limited (Kumar and Leone 1988, Pancras et al. 2012), managers often resort to the strategy of opening new outlets to achieve higher sales. However, the effect of opening new outlets on the performance of a whole retail chain is complex (Pancras et al. 2012). Picking the outlet that generates the largest expected sales on its own will not guarantee the largest marginal revenue of the whole retail chain.

Location is important to retailers, especially for those whose services have to be made through fixed geographic locations (Brown et al. 1981, Ketelaar et al. 2017). Access to retail outlets is a critical factor in determining patronage (Eze et al. 2015, Ghosh and Craig 1986). A good location strategy also gives the retailer strategic advantages over its competitors (Jain and Mahajan 1979, Obeng et al. 2016). The location decision usually represents a large-amount long-run fixed investment (Craig 1984, Öner 2016). When considering adding more service locations, the retailer trades off between the expected revenue with the additional cost of providing the service (Ghosh and Craig 1986, Pancras et al. 2012). In other words, there shouldn't be expansion unless there is sufficient increase in profit. Moreover, the manager should not only consider a single separate location, but also the whole retail network.

This paper focuses on retail network expansion in the context of *add-on products*. The specific products that we study in this paper are those for which:

1. demand is realized only as a supplement to the demand for another product (which we refer to as the *base product*)
2. purchases are made as optional additional purchases at the point of sale of base products
3. base-product sales are done in standard brick-and-mortar (non-online) retail outlets.

The add-on products we study are therefore of the following form. First, the demand is derived from the demand for another product. Second, the add-on product can only be sold together with the base product. Third, the base product is sold at certain physical locations, and both products are purchased at these locations.

There are many real-world settings where products of this form are sold. One such setting is in *car maintenance*. For example:

1. Fuel additives sold at-the-pump; the base product is gasoline, and the add-on product is the additive.
2. Tire polish services sold at car wash places; the base service is a car wash, and the add-on product is the tire polish.

3. Air filter replacement at oil change; the base product is the oil change, and the add-on product is the air filter.

The company selling the add-on product that we partner with is an add-on retailer for a car maintenance product, and so we focus the discussion on this context. Note that the setting is not limited to car maintenance, where similar examples are found ubiquitously in many industries, for example with body care products sold during personal care services and private training lessons in fitness clubs, for example.

The problem we study is how to maximize expected sales of add-on products by expansion, where a company can only expand to a limited number of locations among those where the base product is currently offered. We present a decision-making framework for determining the optimal location for expansion. This requires accurate predictive models of demand given a set of base product retail locations, and an optimization framework for automating the selection of an optimal collection of expansion sites.

For predicting sales, we apply several machine learning algorithms and evaluate their performance on a year of point-on-sales data at existing sites where add-on products are sold. We use basic area-specific demographic information (median household income and population) to supplement base-product and add-on product demand data to create more accurate models.

We further supplement the data set with a spatial-weight matrix, whose entries provide a measure of geographic closeness between base-product retail locations that decays as the distance grows. Including this matrix as part of an input to the learning models is shown to increase predictive accuracy only when spatial autocorrelation exists in base-product demand, as exhibited through an experimental evaluation on two different regions, one in the presence of and the other in the absence of spatial autocorrelation, as indicated by the *Moran's I Test* (Bivand et al. 2008, Moran 1950).

The models we develop in this paper have predictive power better than other models suggested in the literature for similar contexts; for example, the (out-of-sample) *mean absolute percent error* of our predictive models is around 20%, which is less than that reported by Fisher et al. (2016), for predicting sales at online websites.

Equipped with an accurate predictive model for sales in a region, we then turn to formulation and solving an optimization model for automating expansion decision making. The optimization model can take various forms, depending on the predictive model employed. We describe in detail the form the optimization model takes for the various predictive models tested, and suggest a simple general-purpose heuristic that can be used for any such model.

In comparison with a standard baseline policy for expansion, in which the add-on product company expands based on higher base-product sales, the predictive-and-optimization modeling framework developed can result in over 5% additional sales, depending on the base-product demand distribution, the number of expansion sites that the add-on product company can afford, and the presence of spatial autocorrelation, simply by choosing better expansion locations. The solutions obtained are shown to be robust, in that they result in favorable expansion decisions, even when evaluated on different predictive models than the one used in the optimization model.

The contributions of this research and the managerial insights they reveal are therefore as follows:

- the design of accurate predictive models for the demand of add-on products using limited data;
- an understanding of when spatial-weight matrices can improve predictive accuracy based on spatial autocorrelation;
- an optimization framework for automating expansion decisions based on predictive models; and
- an application to a real-world data set where the predictive-and-optimization modeling framework provides substantially better solutions than baseline expansion policies.

The remainder of the paper is organized as follows. We first review related literature in Section 2. In Section 3 we provide details of the problem setting and context. In Section 4 we describe the data set provided by our industrial partner, and the various data elements that we gathered from publicly available sources. Section 5 describes the predictive models we employ and analyzes their accuracy on the data. Section 6 presents the optimization models, both in general and for specific predictive models, suggests a general-purpose heuristic for solving the optimization models, and reports on the effectiveness of the solution obtained. We conclude in Section 7.

2 Literature Review

This work is closely related to a variety of research streams. We discuss each, highlighting the most relevant papers found in the literature to the problem we study.

2.1 Machine Learning for Predicting Demand and Optimizing Decision Making

In recent years, there have been a few studies focused on predicting demand for products using advanced machine learning algorithms. [Ferreira et al. \(2015\)](#) test linear regression models and regression trees for the daily sales of an online retailer, where they find that regression trees provide the best prediction. [Cui et al. \(2017\)](#) implement multiple machine learning methods to forecast daily sales for an online apparel retailer. [Fisher et al. \(2016\)](#) propose a combined method to predict demand at potential locations, which utilizes both machine learning methods and econometric techniques.

Our work differentiates and builds upon this literature by studying add-on products, estimating their demand using machine learning algorithms, and then using optimization to determine best operational decisions. We also introduce the idea of performing a robustness check, where one takes the solutions obtained by a predictive-and-prescriptive model and evaluate its performance on other predictive models. This ensures that the solutions obtained are robust and can be trusted by the organization employing them for decision making.

2.2 Add-on Products

There is a limited stream of research focused on understanding how add-on products affect base-product sales. This research includes studies on how add-on products provide vertical differentiation ([Lin 2017](#)), and how they affect consumers' evaluation of a base product ([Bertini et al. 2008](#)). As an up-selling strategy, there is conflicting research on whether add-on products harm consumer welfare ([Gerstner and Hess 1990](#), [Hess and Gerstner 1998](#), [Wilkie et al. 1998a,b](#)). Other studies related to add-on products include pricing strategies for promotional products under up-selling ([Aydin and Ziya 2008](#)) and pricing strategies for firms selling both base products and add-ons ([Erat and Bhaskaran 2012](#)).

To the best knowledge of the authors, this paper provides the first study focused on demand estimation and retail expansion for add-on products. We furthermore distinguish from the literature, as we do not presuppose that the base product and the add-on product belong to the same firm.

2.3 Retail Site Decision Making

Location decision making in the retail industry is critical. There exists considerable literature dealing with retail location models ([Craig 1984](#), [Drezner et al. 2015](#)). The research most closely

related to the work in this paper is by [Ghosh and Craig \(1986\)](#), where the authors design a network of service centers via choosing the network size, the location of outlets, and operating characteristics simultaneously. A major difference between this work and the present paper is that the authors do not consider competition among outlets. In addition, [Ghosh and Craig \(1986\)](#) deals with determining an optimal network design, including both the optimal size of the network given an exogenous demand and the characteristics that will attract customers. In our setting, the network is fixed, in that expansion decision for the add-on product is restricted to the locations that sell the base product.

There is an increasing number of papers that consider the cannibalization effect among outlets of the same retailer ([Fisher et al. 2016](#), [Nishida 2014](#), [Pancras et al. 2012](#)). [Pancras et al. \(2012\)](#) examines the impact of opening/closing an outlet on overall chain performance, but their demand model is more suitable for isolating locations that can be closed. Our focus is narrowed to expansion decision. Due to the add-on nature of our focal product, we operate under the assumption that the cost of keeping operation is lower than that of ending the contract with the institutes that operate the base product. Second, like most research on this topic, they estimate the parameters using panel data from a large U.S. city. This causes context dependency. Results from one study area cannot be expected to hold in other areas. This limitation restricts the use of choice models calibrated from one study region to design location strategies in another.

2.4 Facility Location Problem

The literature on facility location problems dates back at least one-hundred years ([Weber et al. 1909](#)), with perhaps an even longer history dating back to studies of Pierre de Fermat, Evangelista Torricelli (a student of Galileo), and Battista Cavallieri ([Drezner and Hamacher 2002](#)). There are many books dedicated to the subject (e.g., [Church and Murray 2009](#), [Daskin 2013](#), [Drezner and Hamacher 2002](#), [Francis and White 1974](#), [Handler and Mirchandani 1979](#), [Love et al. 1988](#), [Rosenwein 1994](#), [Verter and Murat 2007](#), [Wolf 2011](#)), as well as countless surveys and articles beyond this. Some authors have attempted to develop a taxonomy for facility location problem ([Daskin 2008](#), [Hamacher et al. 1996](#)) where models are broken down based on continuous/discrete/network demand and location decision, the number of facilities, the type of facilities, capacity restrictions, and objective functions.

The stream of literature most-closely related to the problem studied in this paper is the *competitive facility location* (CFL) model where facilities set up by decision makers will compete for market share and profitability. Hotelling ([Hotelling 1929](#)) was perhaps the first to study

location models with competition where he considers the case of two ice cream shops picking locations for their respective storefronts and customers choose a shop based on their relative distances. There have been a number of papers dedicated to classifying CFL models (Drezner and Zemel 1992, Eiselt et al. 1993, Karakitsiou 2015, Kress and Pesch 2012). The taxonomy is typically based on the spatial representation (continuous, discrete, network), distance measure (Manhattan distance, $\ell - 2$ norm, etc.), the nature of the competition (static, dynamic, sequential), the number of new locations (single versus multiple), the nature of the customers (demand is elastic versus inelastic, customers come from a single location or multiple locations, customers are served by a single location or by any number of locations based on a probability distribution, etc.), the objective type (customers prefer facilities nearby or far way), among other defining characteristics.

CFL models using spatial interaction typically employ a demand function for each customer, and that model is based on the probability that a customer will choose a particular facility based on the customer’s proximity to a facility. In our model, customer demand for our focal products is a function of the base-product sales and the sales in nearby locations (scaled by their proximity).

Studies do appear in the literature that investigates CFL with cannibalization and market expansion. Bozkaya et al. (2010) study CFL with location routing where a firm incurs costs due to vehicles having to service each location from a central warehouse. Other authors have investigated spatial interaction models using concave demand models (Aboolian et al. 2007a,b, 2009) but again do not consider demand as derived from other products. More recently there have been papers investigating models using a leader-follow framework, where one firm decides where to locate facilities and the other firm follows with a decision (Drezner et al. 2015).

The problem studied in this paper differs from those appearing in the literature in that we study the demand of our focal products as a derived demand for another product.

3 Problem Context

The goal of this paper is to develop a hybrid predictive-and-optimization modeling framework for automating the retail expansion location decision-making process for companies selling add-on products where base products are sold in fixed geographic locations. We, therefore, seek to develop predictive models specifically designed for derived-demand products, and identify expansion locations based on those predictive models. We first develop a relationship between the profitability of retail outlets specifically for our focal product and the characteristics of their

locations, and then we optimize the profit of the entire network.

We learn the relationship between the performance and the location characteristics from studying the past performance of existing locations. When retail managers make location expansion decisions, the attractiveness level of a location includes the demand, the buying power, and the competition level are common factors to consider. In addition, the demand for our focal product is derived from the demand for the base product. Hence, in the predictive models of the focal product’s sales, we use the base-product sales, the weighted base-product sales in the neighborhood, the local median household income, and the local population as our predictive variables.

Before proceeding with the models, we fix notation. Let $N = \{1, \dots, n\}$ be the set of sites selling the base product. N is partitioned into two sets, S and Q , representing the sites currently selling the add-on product (*active* sites) and the possible sites for expansion (*candidate* sites), respectively. Denote by $n := |N|$, $s := |S|$, and $q := |Q|$ the cardinality of the sets. Let g_i be the total sales of the base product at site i . Furthermore, let K be the (maximum) number of expansion sites that the firm can choose. We assume that expanding operations to any site requires a fixed cost which does not vary from site-to-site, and so, given a fixed budget for expansion, the firm can choose any set of sites within the allowable budget, or through management choice, to expand to, although this assumption can be relaxed.

In our setting, an add-on product is offered for purchase each time a customer purchases the base product. For any site $i \in N$, let g_i be the number of times the base product is purchased at site i (i.e., demand for the base product at size i). In order to build accurate models for predicting purchases, we use demographic information common to the field studied. In particular, we incorporate median household income h_i and zip-code population p_i . For notational convenience, we drop the index on parameters to represent the entire vector of values of that parameter (e.g., $g \in \mathbb{R}^n$ is the n -dimensional vector of demands for all sites).

The first step in the framework is to develop a robust and accurate predictive model for the sales of the add-on product, a . More formally, for any $\tilde{N} \subseteq N$, let $f(\tilde{N})$ be the amount of the add-on-product sales given the firm’s selection to offer the add-on product at locations \tilde{N} . We seek a model \hat{f} that approximates f , using known data, which includes base-product sales at each location and median household income and population in the area of each location.

Our models allow for the incorporation of location effects. In particular, the data we tested suggests that even for the same add-on product, that location effects can exist in particular regions or not. This can be formally tested through checking for *spatial autocorrelation* through

the Moran’s I test.

After determining the best predictive model, we formulate a general optimization model for finding the best expansion sites. Namely, we solve:

$$\max_{\tilde{N} \subseteq N} \left\{ \hat{f}(\tilde{N}) : \text{subject to } \tilde{N} \supset Q, |\tilde{N}| = |\tilde{S}| + K. \right\} \quad (\text{EXP})$$

The complexity of this optimization problem depends on the underlying structure of \hat{f} . We discuss the various forms this model takes based on the predictive model employed in Section 6. We also devise a simple greedy heuristics for solving (EXP) given any functional form, which, despite its simplicity, works well and finds optimal solutions for all cases tested. For the problems in which we can find optimal solutions, the predictive models are of a form that when formulating (EXP), the problem reduces to a simple linear or quadratic binary optimization problem.

4 Data

In order to evaluate the efficiency of the framework we develop, we use data provided by an add-on product retailer. The retailer operates in over 1500 locations in the US, and partners with seven base-product companies, operating in locations that they currently do and can in the future sell their products. The specific product is an add-on product in the context of car maintenance, that is sold at locations where customers of the base-product frequent and are offered the product on site at the time of purchase.

The data utilized is at-the-site transactional data, that indicates each time a customer decides to purchase the add-on product when purchasing the base product. This option is available each time the customer purchases.

We select two geographical regions in the U.S. to study. Region 1 covers about 11878.59 *sq mi*, and has 89 existing operation sites and 228 candidate sites. The average distance between locations in this area is 17.2 miles. The locations are closely gathered as one cluster. Region 2 covers about 132549.2 *sq mi*, and has 146 existing operation sites and 162 candidate sites. The average distance between locations in this area is 54.8 miles. There are three clusters in this area. Figure 1 depicts these regions. In both figures, the circles with crosses represent current locations, i.e., locations in S . The hollow circles represent candidate locations, i.e., locations in Q . We can see from Figure 1 that there are relatively more expansion choices in Region 1 compared to Region 2. In both regions, the candidate sites are not evenly located. Some of

them are more closely clustered, while some are far from any center. Our expansion strategy works well for both regions as demonstrated in Section 6.

For each site $i \in S$, we have the aggregated point-of-sale transaction records of the outlets during 2015. We have the transaction records of both the base product and the add-on product. In addition, we have their geographic location (the latitude and the longitude in degrees), and the demographics, i.e., the median income per household and the population in a corresponding zip code. In summary, for each outlet, we have the number of the base product annual transactions, the number of the focal product annual transactions, its latitude and its longitude in degrees, the median income per household and the population in the corresponding zip code.

The candidate locations Q we consider are those locations run by the base-product retailer where the add-on product is not currently sold. By the same mechanism used to gather data about current add-on sites, we have, for each location $i \in Q$, h_i and p_i . The data from our partner company only lists base-product sales for the sites that the add-on product is currently sold at. This provides us an opportunity to test different base-product demand profiles, in order to understand when the models we propose are particularly effective. We discuss this in Section 6.

Summary statistics of the data is provided in Table 1. For Region 1 and Region 2, for S and Q , the data reports the mean and standard deviation of the base-product sales g , the add-on product sales a , the pair-wise distance between locations D (in miles), the household income h and the population p .

Table 1: Summary statistics of industry partner data

	Region 1				Region 2			
	S		Q		S		Q	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
g	196700.20	82871.41	—	—	195246.30	93927.96	—	—
a	3420.58	1149.45	—	—	3087.76	1548.83	—	—
D	34.67	18.21	31.37	18.30	163.81	88.08	150.61	98.03
h	67240.60	18033.64	62505.84	11890.57	62434.31	22703.08	69004.53	18194.45
p	50641.40	17718.66	45898.06	12815.12	47208.74	24986.17	53736.55	15935.17

5 Predictive Models

This section provides details of the predictive models applied which will be used to optimize expansion decisions.

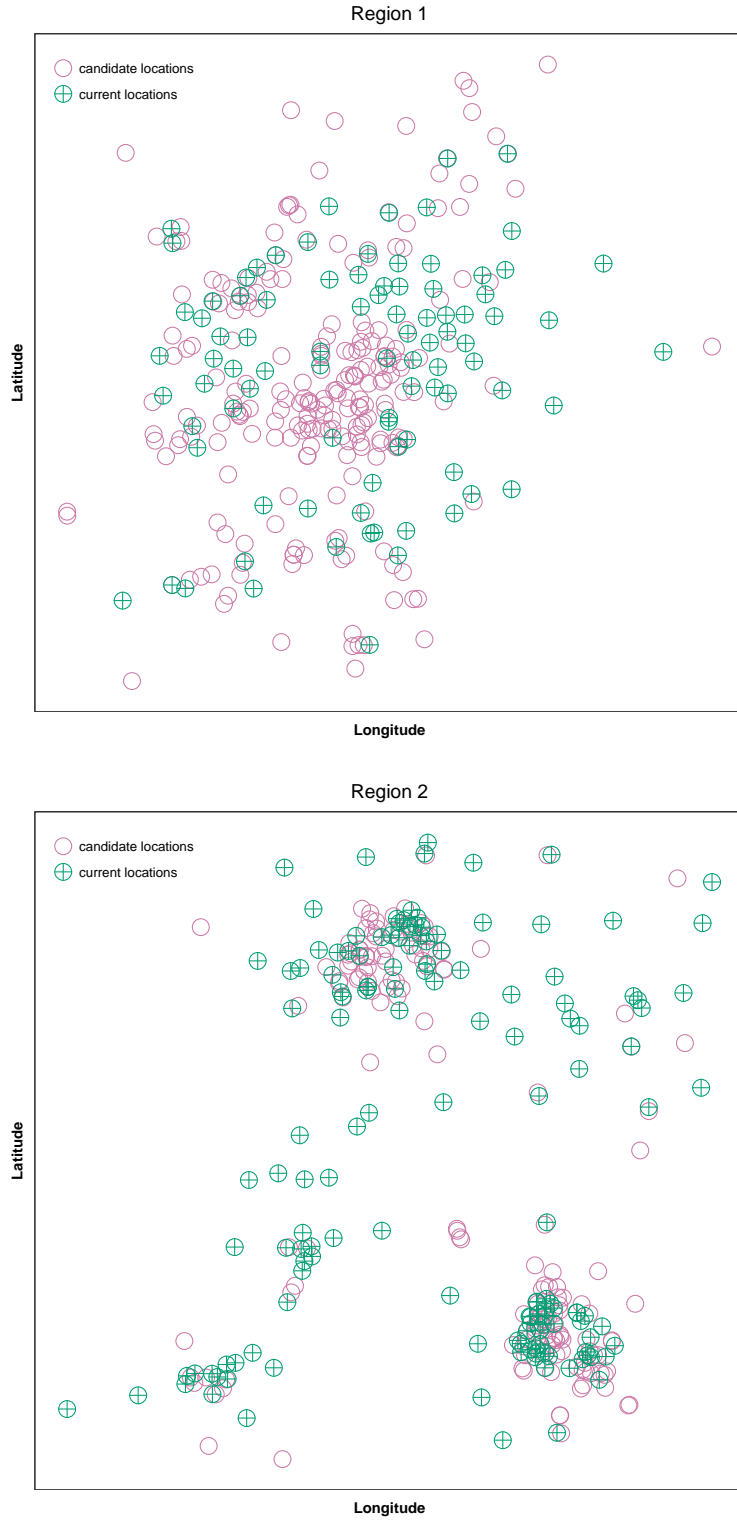


Figure 1: Spatial distribution of base-product retail sites in Region 1 and Region 2.

5.1 Testing for Spatial Autocorrelation

Before building the predictive models, we first look into the spatial interaction of both base-product sales and add-on-product sales, in the current sites. First, georeferenced observations, for example, the two product sales at locations in vicinity, generally are not independent of one another (Getis 2008). Strong evidence of the influence of spatial autocorrelation in the consumption of gasoline has been found (Dos Santos et al. 2014). Second, whether or not there is an impact on sales (either positive or negative) of either based on proximity can impact the best choice of predictive model. We find that the best predictive model for a given region depends on whether or not spatial autocorrelation in the base-product sales exists. As we will show, using or omitting a feature representing the total base-product sales scaled by the distance from the focal location has a positive impact on predictive models in the presence of spatial autocorrelation, and a negative impact in its absence.

In order to test whether there exists spatial autocorrelation in base-product sales, we apply a *Moran's I* test. This test is specifically designed to test whether or not spatial autocorrelation exists. The test statistics in this test, I_g , is calculated as a ratio of the product of the variable of interest and its spatial lag, with the cross-product of the variable of interest, and adjusted for the spatial weights used:

$$I_g = \frac{s}{\sum_{i \in S} \sum_{j \in S} w_{i,j}} \cdot \frac{\sum_{i \in S} \sum_{j \in S} w_{i,j} (g_i - \bar{g})(g_j - \bar{g})}{\sum_{i \in S} (g_i - \bar{g})^2},$$

where \bar{g} is the average base-product sales over sites in S . The test assumes, in the absence of spatial autocorrelation, that this statistic follows a normal distribution. In other words, if the test result is not significant, we cannot reject the null hypothesis that the observations are random. We discuss the results of applying the test to the sites in both regions in Section 5.3.

The feature we include to encode the scaled base-product sales is an inverse distance matrix W times g , the vector of gas sales. In particular, for any two sites $i, j \in N$, define $w_{i,j} := \frac{1}{d_{i,j}}$, where $d_{i,j}$ is the geographic distance (taken as the Euclidian distance) between the two sites. In this way, Wg is a vector that in coordinate i contains the value $\sum_{j \in \tilde{N}, j \neq i} \frac{1}{d_{i,j}} g_j$.

5.2 Demand Prediction Models

In order to create the most accurate predictive models for estimating base-product sales, we apply linear regression models as a parsimonious baseline, and *support vector regression* (SVR) models as a more advanced and nonparametric method for predicting expected add-on product

sales. Our selection of predictive models was done for the following reasons. Linear regressions is a classical tool that has been used ubiquitously throughout the literature, and provide a good baseline. SVRs are used because they provide another more advanced mechanism for predicting sales, which allows input data to be transformed via *kernel* functions in order to accommodate nonlinearity. Our choice of prediction model among other available options (e.g., random forests, neural networks) is that SVRs provide closed-form solutions (which is critical for optimal expansion decision making), and that in other papers estimating demand using predictive models, SVRs provide among the best predictive accuracy (Cui et al. 2017, Fisher et al. 2016). Random forests are shown to provide outstanding predictions (Cui et al. 2017, Fisher et al. 2016). However, random forests do not work well in our spatial setting with a limited set of features used for prediction. Random forests rely on resampling both features and data in order to build prediction models. Features we employ cannot be expected to offer significantly more information when being resampled in building the random forests. Second, the construction of W relies on the entire training data. When data is resampled with replacement, it is unreasonable to have two sites at the same location. Consequently, random forests would be nearly identical to a single decision tree if we resample without replacement. Therefore, we employ SVR models and linear regression models for prediction. In our preliminary tests, the linear-kernel and radial-kernel SVR models outperform those with a sigmoid kernel or a polynomial kernel, and therefore we only apply these for the remainder of this paper.

Our goal is to predict the add-on product sales given any set of sites. In our derived demand scenario, the demand for the add-on product is derived from the demand for the base product. Hence, a portion of the predicted add-on product sales comes from sales of the base product. In addition, the probability that a customer patronizes a facility is proportional to the attractiveness level and to a distance decay function (Berman et al. 2009) and the attractiveness of a location is related to its demographics. We include the base-product sales, the local median income per household and the total population as the independent variables, along with Wg .

Price might be a strong indicator for sales. However, in our case, the price of the add-on product is fixed across the year we study. The price of the base product doesn't vary much across locations, especially after we take the average of the entire year. The mean of the base-product price in Region 1 is 2.21, while its standard deviation is 0.076. The mean of the base-product price in Region 2 is 2.09, while its standard deviation is 0.071. Hence, we don't include the price in our prediction models.

In summary, the predictive models we apply are as follows (In regions without the spatial

autocorrelated base-product sales, we remove Wg):

$$\begin{aligned} \text{Linear Regression:} & \quad a = \beta_1 g + \beta_2 Wg + \beta_3 h + \beta_4 p + \beta_0 + \epsilon \\ \text{Linear-Kernel SVR:} & \quad a = \omega \cdot \ker_{LK}(g, Wg, h, p) + b_0 + \epsilon \\ \text{Radial-Kernel SVR:} & \quad a = \omega \cdot \ker_{RK}(g, Wg, h, p) + b_0 + \epsilon \end{aligned}$$

where,

- a is the add-on product sales;
- g is the base-product sales;
- Wg is the spatially weighted average of the base-product sales;
- h is the local median household income;
- p is the local population;
- ω is the estimated weights in the SVR models;
- \ker_{LK} is the linear kernel function, i.e., (x, x') ;
- \ker_{RK} is the radial kernel function, i.e., $\exp(-\gamma\|x - x'\|^2)$;
- b_0 is the estimated constant term from the SVR models; and
- ϵ is the error.

5.3 Analysis

We begin with the results from the Moran's I tests, obtained using the `moran.test` (for the sales of the two products) and `lm.morantest` (for the residuals of a linear regression model of the add-on-product sales on the base-product sales) (in package `spdep` in R, version R.3.4.2) (Bivand et al. 2008). Table 2 reports the p -values of the Moran's I tests for the base-product sales, the add-on-product sales, and the residuals, in both regions of interest. Neither tests of the two product sales result in statistically significant spatial autocorrelation for Region 1, while with the opposite true for Region 2. We, therefore, conclude the following:

- Neither the base-product sales nor the add-on-product sales in Region 1 are spatially autocorrelated.
- Both the base-product sales and the add-on-product sales in Region 2 show spatial autocorrelation. In addition, after controlling for base-product sales, the residuals of the add-on-product sales show insignificant spatial autocorrelation, allowing us to conclude that the spatial autocorrelation in the focal product sales is inherited from the spatial autocorrelation of the base-product sales.

Table 2: Results (p -values) from Moran’s I test

Region	Add-on-product sales	Base-product sales	Residuals
Region 1	0.2568	0.1986	0.6506
Region 2	0.02739**	0.05601*	0.6291

Generally speaking, if the Moran’s I test statistic is statistically significant and positive, then the spatial distribution of high/low values is more spatially clustered than would be expected if underlying spatial processes were random (ArcGIS 2017). Based on the test results in Table 2, this is the case for Region 2. We can observe three clusters on Figure 1. On the other hand, for Region 1, we cannot reject the possibility that the spatial distribution of feature values is the result of random spatial processes.

We conduct grid search over 50-time 10-fold cross validation to tune the hyper-parameters in the two SVR models. For the linear-kernel SVR model, we need to tune the cost C and ε . For the radial-kernel SVR model, we need to tune C , γ , and ε . We compute the *root mean squared error* (RMSE) as a measurement of how the combination of the hyper-parameter values performs. We pick the setting of the hyper-parameter values with the lowest RMSE. We start search with C in $\{2^0, 2^1, \dots, 2^{16}\}$, $\varepsilon \in \{0, 0.1, \dots, 1\}$, and $\gamma \in \{10^{-7}, 10^{-6}, \dots, 10^{-3}\}$. If the optimal combination contains boundaries, then we adjust the search range to make sure the previous optimal hyper-parameter value is within the search range. We iterate this procedure until the optimal hyper-parameters found are inside the search ranges. We report the value of the hyper-parameters we pick in Table 3.

Table 3: Best hyper-parameters

Region	Model	C	ε	γ
Region 1	Linear-Kernel SVR (4 features)	0.0625	0.3	—
Region 1	Radial-Kernel SVR (4 features)	2^{12}	0.3	10^{-5}
Region 1	Linear-Kernel SVR (3 features)	2^{14}	0.5	—
Region 1	Radial-Kernel SVR (3 features)	2^{10}	0.5	10^{-3}
Region 2	Linear-Kernel SVR (4 features)	2^8	0.1	—
Region 2	Radial-Kernel SVR (4 features)	2^{16}	0.3	10^{-4}
Region 2	Linear-Kernel SVR (3 features)	2^6	0.1	—
Region 2	Radial-Kernel SVR (3 features)	2^{10}	0.2	10^{-3}

We now turn to evaluating the impact of including Wg as an input variable in the models tested. As described below, we find that in the region exhibiting spatial autocorrelation, it’s beneficial to have the spatial weighted average base-product sales as one of the features.

Table 4 and 5 reports the *root mean squared error* (RMSE) and the *mean absolute percent error* (MAPE) based on a 50-time 10-fold cross validation for Region 1 and Region 2, respectively. For Region 1, which exhibits no spatial autocorrelation, models using three features are significantly better than those with four features. Radial-kernel SVR models slightly outperform the other two types of models. The out-of-sample MAPE of this radial-kernel SVR model is 22.9%, which is lower than those in Fisher et al. (2016). For Region 2, which has spatial autocorrelation, models using four features are slightly better than those with three features. Similar to Region 1, radial-kernel SVR models outperform the rest two. Also, the out-of-sample MAPE of this model is 16.8%, which is lower than that those in Fisher et al. (2016). We therefore conclude that radial-kernel SVR models are best for both regions, and that including Wg for Region 2 increases predictive accuracy.

Table 4: Predictive Models Performance of Region 1

Region 1	Linear-Kernel SVR	Radial-Kernel SVR	Linear Regression
Mean RMSE (4 Features)	751.40	748.99	835.75
SD RMSE (4 Features)	14.73	16.20	18.85
Mean MAPE (4 Features)	24.51%	24.19%	23.35%
SD MAPE (4 Features)	0.80%	0.91%	1.08%
Mean RMSE (3 Features)	719.99	718.78	733.80
SD RMSE (3 Features)	9.28	9.61	10.36
Mean MAPE (3 Features)	23.33%	22.90%	23.67%
SD MAPE (3 Features)	0.29%	0.34%	0.40%

Table 5: Predictive Models Performance of Region 2

Region 2	Linear-Kernel SVR	Radial-Kernel SVR	Linear Regression
Mean RMSE (4 Features)	771.38	756.47	788.43
SD RMSE (4 Features)	4.82	8.40	8.56
Mean MAPE (4 Features)	17.06%	16.84%	18.02%
SD MAPE (4 Features)	0.31%	0.27%	0.40%
Mean RMSE (3 Features)	772.75	762.03	791.83
SD RMSE (3 Features)	2.82	6.40	6.60
Mean MAPE (3 Features)	16.37%	16.26%	17.92%
SD MAPE (3 Features)	0.16%	0.22%	0.25%

In the optimization models we use for determining expansion decisions, we will fix the best predictive models for each region and find expansion sites that will maximize expected sales. The important insight from the analysis conducted in this section is that in the presence of spatial autocorrelation, including the weighted matrix is critical for good predictive performance.

6 Expansion Optimization

Equipped with a high-performing predictive model, we now turn to determining the optimal set of expansion sites, to maximize expected sales for add-on products given the set of candidate locations. The form of the optimization model depends entirely on the choice of predictive models. Simpler predictive models result in tractable optimization problems, and more complex models result in highly nonlinear optimization models that require heuristics.

As described in Section 3, the optimization problem, in its most general form, is (EXP). This section first explores how this optimization model is specified for the various predictive models employed, and then presents a general purpose greedy heuristic for finding expansion decisions. We then describe how the solutions we obtain perform in practice on the data made available from our industry partners.

6.1 Models with No Spatial Effects

Consider a predictive model \hat{f} taking the form

$$\hat{f}(\tilde{N}) = \sum_{i \in \tilde{N}} l_i. \quad (\text{LM-NS})$$

This results from, for example, if a linear regression model or a linear-kernel SVR model is employed. In the former, $l_i = \beta_1 g_i + \beta_3 h_i + \beta_4 p_i + \beta_0$. In the latter, $l_i = \sum_{k=1}^m C_k (V_{k,1} g_i + V_{k,3} h_i + V_{k,4} p_i) + b_0$, where C_k and $V_{k,j}$ ($j = 1, 2$, and 4 in our example) are the standard parameters in SVR models, and m is the number of support vectors of a specific SVR model. In either event, (EXP) reduces to

$$\begin{aligned} \max \quad & \sum_{i \in N} l_i x_i \\ \text{s.t.} \quad & \sum_{i \in Q} x_i = K \\ & x_i = 1, \quad \forall i \in S \\ & x_i \in \{0, 1\} \quad \forall i \in Q. \end{aligned}$$

Notice that this expansion optimization problem does not involve any association with existing sites that the add-on product is sold at. Furthermore, this can be solved by simply sorting the predicted add-on-product sales at the candidate sites in a non-increasing order of l_i , and choosing the first K sites to expand to.

6.2 Linear Models with Spatial Effects

Suppose we incorporated the weight matrix W in a linear regression or linear-kernel SVR. Given a collection of sites \tilde{N} offering the add-on product, the model will predict

$$\hat{f}(\tilde{N}) = \sum_{i \in \tilde{N}} l_i x_i + \sum_{i \in \tilde{N}} \sum_{j \in \tilde{N}, j \neq i} e_{i,j} x_i x_j \quad (\text{LM-WS})$$

as the total sales of the add-on product. For a linear regression model, we have

$$l_i = \beta_1 g_i + \beta_3 h_i + \beta_4 p_i + \beta_0, e_{i,j} = \beta_2 \sum_{j \neq i} \frac{1}{d_{i,j}} g_j,$$

and for a linear-kernel SVR, we have

$$l_i = \sum_{k=1}^m C_k (V_{k,1} g_i + V_{k,3} h_i + V_{k,4} p_i) + b_0, e_{i,j} = \sum_{k=1}^m C_k (V_{k,2} (\sum_{j \neq i} \frac{1}{d_{i,j}} g_j)).$$

The expansion optimization model can then be written as

$$\begin{aligned} \max \quad & \sum_{i \in N} l_i x_i + \sum_{i \in N} \sum_{j \in N, j \neq i} e_{i,j} x_i x_j \\ \text{s.t.} \quad & \sum_{i \in Q} x_i = K \\ & x_i = 1, \quad \forall i \in S \\ & x_i \in \{0, 1\} \quad \forall i \in Q. \end{aligned}$$

The ensuing model is, therefore, a *cardinality-constrained binary quadratic* optimization problem (Bertsimas and Shioda 2009), which has large literature containing dedicated methods. For this paper and the data set we use, the optimization model was easily solved by a commercial integer programming solver (GUROBI 7.5.1).

6.3 More General Models

Nonlinear predictive models often result in superior predictive power, as in the case for Region 2. The resulting optimization models are more complicated. For example, consider a radial-kernel SVR model with a spatial weight matrix as a part of a feature. Given a collection of sites \tilde{N} , the function form is more complicated than the linear or quadratic form above. We, therefore, devise a heuristic to solve the ensuing optimization model.

Given any function $\hat{f}(\tilde{N})$ and a number of expansion sites K , one can devise a host of heuristics to plan expansion. We suggest the following, simple heuristic, that can apply to any analytical function.

Start with no candidate sites, and let $\tilde{N} = S$. Among all sites in Q , choose the site Φ that, if added to \tilde{N} , the increase from $\hat{f}(\tilde{N})$ to $\hat{f}(\{\tilde{N}, \Phi\})$ is the most. Then, remove Φ from Q , and add it to \tilde{N} . Continue until $|\tilde{N}| - |S| = K$.

This heuristic does not guarantee an optimal selection of sites. It will guarantee an optimal selection when employing models with no spatial weight matrix. Also, in our experimental results, this heuristic was able to match the optimal expansions for all linear models with spatial weight matrix. There is no guarantee that this will happen for other data sets, but given its performance, we use this simple heuristic for the more complicated nonlinear radial-kernel SVR model as well.

6.4 Analysis of Results

In order to evaluate the performance of the expansion optimization framework, we compare with a standard practice that an add-on retailer might employ, which determines expansion by choosing those sites for which the base-product sales are highest. We refer to this expansion as the *baseline* algorithm, or **BM**. For any predictive model P , we let **EO-P** be the expansion optimization algorithm using predictive model P .

In order to simulate base-product sales in candidate locations, for each $i \in Q$, we draw, independently, g_i from a normal distribution with mean μ and standard deviation σ . μ is taken as the mean of base-product sales in S , and we test for $\sigma = 4^s$, for $s \in \{2, 4, 6\}$, in order to evaluate differences in performance for different base-product sales distributions. If g_i is drawn to be below the minimum base-product sales in the region, we set g_i equal to that minimum. We simulate this demand ten times per σ for the following experiments. We test for $K \in \{1, 2, \dots, 20\}$.

We first compare, for the best predictive model in each region, P^* , what the average predicted total sales of add-on product will be, using both **BM** and **EO-P***, for each K tested. P^* represents radial-kernel SVR models for both regions, with Wg included as an independent variables in Region 2.

Table 6 report results comparing the expansion decisions. Each row corresponds to varying K . The columns report the average increase in additional sales that result from employing **EO-P*** over **BM**. In particular, let z^0 be the total sales of add-on product estimated by P^*

Table 6: Average percent increase in additional sales by using **EO-P*** over **BM** in both regions for various base-product demand distributions

K	Region 1			Region 2		
	$s = 2$	$s = 4$	$s = 6$	$s = 2$	$s = 4$	$s = 6$
1	5.48%	4.45%	2.65%	3.18%	3.45%	2.11%
2	4.58%	4.39%	2.65%	4.43%	4.21%	2.26%
3	4.64%	4.22%	2.47%	5.50%	4.54%	2.43%
4	4.52%	4.03%	2.25%	5.63%	5.64%	2.67%
5	4.59%	4.07%	2.17%	5.51%	5.54%	2.54%
6	4.42%	3.95%	2.07%	5.19%	5.57%	2.64%
7	4.30%	3.78%	1.93%	4.88%	5.98%	2.58%
8	4.21%	3.66%	1.84%	4.69%	5.67%	2.43%
9	4.18%	3.65%	1.77%	4.55%	5.46%	2.62%
10	4.00%	3.57%	1.75%	4.44%	5.32%	2.78%
11	3.89%	3.51%	1.70%	4.38%	5.10%	3.03%
12	3.85%	3.49%	1.63%	4.50%	5.08%	2.96%
13	3.79%	3.44%	1.59%	4.68%	4.96%	2.84%
14	3.72%	3.43%	1.51%	4.53%	4.79%	2.86%
15	3.61%	3.34%	1.52%	4.39%	4.60%	2.96%
16	3.54%	3.29%	1.51%	4.28%	4.45%	2.95%
17	3.49%	3.21%	1.48%	4.27%	4.35%	3.02%
18	3.44%	3.17%	1.44%	4.22%	4.43%	2.91%
19	3.41%	3.13%	1.43%	4.33%	4.54%	2.95%
20	3.38%	3.10%	1.39%	4.62%	4.64%	2.94%

without any expansion. For any K , let z_e^K be the total add-on sales predicted using **EO-P*** and z_b^K be the total add-on sales predicted using **BM**. The table reports $\frac{z_e^K - z_b^K}{z_b^K - z_0^K} \cdot 100\%$. The higher this value, the more gains that are realized by using the framework we propose.

6.5 Managerial Insights

Our results provide clear evidence that making expansion decisions for add-on products should not be done without more readily understanding how sales will be affected by the expansion decisions made. A standard approach to expanding by only considering base-product sales can lead to far worse expansion decisions than the more complex but effective predictive and optimization framework developed in this paper.

As shown in Table 6, when K is small, we realize larger gains when there is no spatial autocorrelation (Region 1). However, this switches when K increases. This indicates that our suggested approach is particularly useful when the retailer is considering a large expansion strategy in the presence of spatial autocorrelation. Also, in general, gains decrease with K when

there is no spatial autocorrelation (Region 1), and increase with K where there is (Region 2).

Table 6 exhibits the gains the add-on product seller can realize by employing our strategy. When comparing with the results obtained with **BM**, the add-on product seller can realize additional sales of over 5%, simply by choosing better expansion decision making. This depends, however, on the distribution of the demand for base-products. For example, if $K = 10$, in Region 1, when the variance is low ($s = 2$), we expect to achieve 4.00% increase in sales simply by choosing a better set of locations. As the variance increases, the difference between **BM** and **EO-P*** becomes smaller. The same trend appears in Region 2 as well.

This trend appears to be more extreme when base-product sales at candidate sites do not vary much among the candidate locations. In such a case, relying on base-product sales does not provide enough information. Hence, the information from the neighborhood become more important. Even when the variation does exist among candidate sites, using simple expansion model relying only on base-product sales will not achieve optimal results. In order to capture complex relationships between demographic information and the interplay between the demand from neighboring base-product retail sites, predictive models using spatial weight matrices can improve prediction power, thereby allowing optimization models to automate the selection of desirable expansion sites.

As far as the net increase in sales, if $K = 20$ and $s = 2$, the difference in additional sales of add-on product in Region 1 is 3429, and in Region 2 it is 3678. For a tight-margin business, this may have significant impacts on its bottom line, particularly because this increase in sales does not require anything more than selecting better expansion sites.

It is also interesting to view the geographic dispersion and selection of sites, comparing the expansion decisions made using **BM** with **EO-P*** for both regions. This is depicted in Figure 2 (for Region 1) and Figure 3 (for Region 2) for one instance generated with $s = 6$. Each figure displays a point for every site at its geographic location. The current sites are squares and colored orange, and the candidate sites are circles, colored blue if chosen for expansion and gray otherwise. Furthermore, the points corresponding to the candidate sites are gradient sized to depict relative base-product sales (higher sales, larger circle). In both figures, the left plot depicts the solution obtained using **BM**, and the right plot depicts the solution obtained using **EO-P***.

These figures depict more clearly how the expansion decisions differ. In particular, we see that the expansion decisions made through **EO-P*** take advantage of local clusters, and although selection leans towards those sites with large base-product sales, it is often better to consider

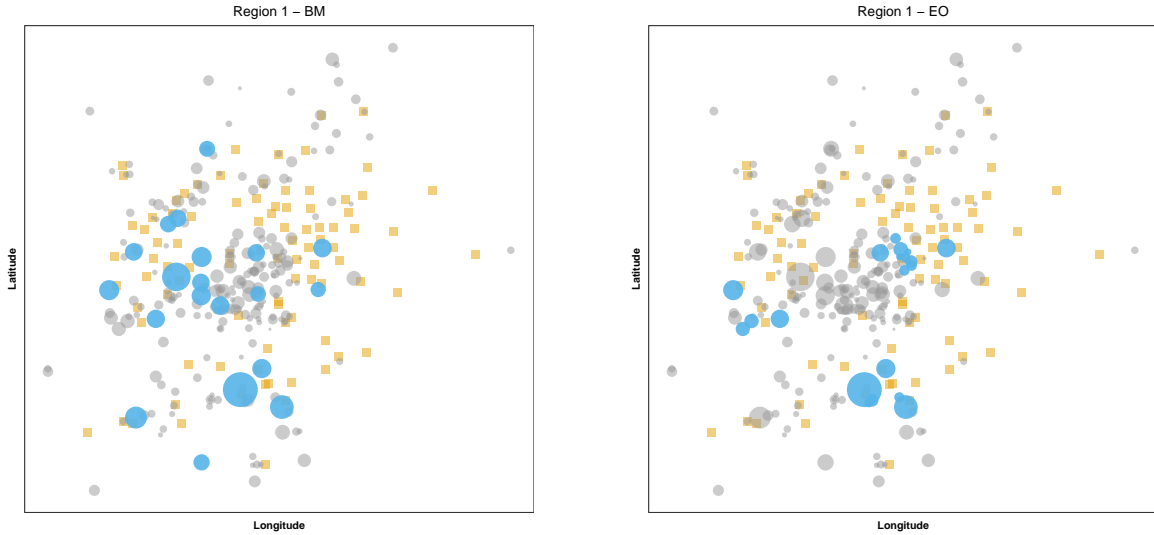


Figure 2: Best known expansion sites for Region 1 using (left) **BM** and (right) **EO-P***. Current sites drawn as yellow squares. Candidate sites drawn as circles, colored blue if selected and gray otherwise. The relative sizes of the circles correspond to the total base-product sales at the corresponding site.

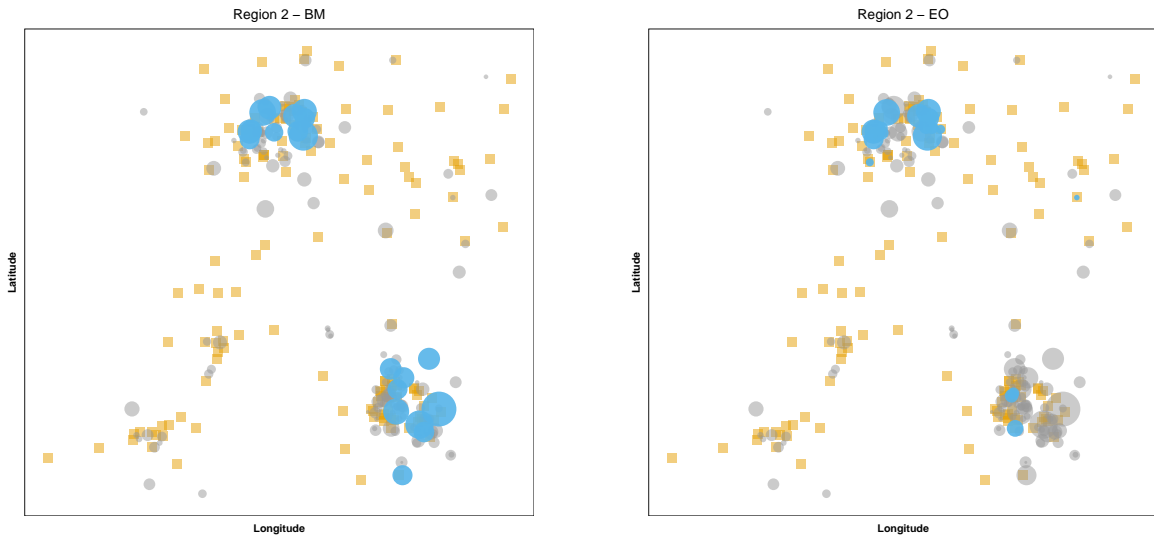


Figure 3: Best known expansion sites for Region 2 using (left) **BM** and (right) **EO-P***. Current sites drawn as yellow squares. Candidate sites drawn as circles, colored blue if selected and gray otherwise. The relative sizes of the circles correspond to the total base-product sales at the corresponding site.

location impact over simply base-product sales.

It is clear that using the optimization models for a fixed predictive model will always result in better results than using the baseline algorithm when evaluated on the predictive model

Table 7: Robustness check for solution obtained by P^* for Region 1

K	$s = 2$			$s = 4$			$s = 6$		
	LR	LK	RK	LR	LK	RK	LR	LK	RK
1	3.34%	4.92%	5.48%	2.61%	3.94%	4.45%	0.61%	1.98%	2.65%
2	2.89%	4.27%	4.58%	2.69%	4.04%	4.39%	0.90%	2.27%	2.65%
3	2.97%	4.39%	4.64%	2.65%	4.03%	4.22%	1.01%	2.21%	2.47%
4	2.92%	4.33%	4.52%	2.53%	3.85%	4.03%	0.90%	2.08%	2.25%
5	2.98%	4.43%	4.59%	2.56%	3.90%	4.07%	0.85%	2.04%	2.17%
6	2.96%	4.43%	4.42%	2.57%	3.94%	3.95%	0.81%	1.96%	2.07%
7	2.90%	4.35%	4.30%	2.46%	3.78%	3.78%	0.73%	1.82%	1.93%
8	2.86%	4.29%	4.21%	2.41%	3.71%	3.66%	0.72%	1.81%	1.84%
9	2.88%	4.33%	4.18%	2.43%	3.75%	3.65%	0.70%	1.76%	1.77%
10	2.79%	4.22%	4.00%	2.40%	3.70%	3.57%	0.68%	1.73%	1.75%
11	2.71%	4.08%	3.89%	2.37%	3.65%	3.51%	0.66%	1.66%	1.70%
12	2.67%	4.02%	3.85%	2.36%	3.64%	3.49%	0.64%	1.61%	1.63%
13	2.62%	3.94%	3.79%	2.33%	3.58%	3.44%	0.63%	1.55%	1.59%
14	2.56%	3.84%	3.72%	2.31%	3.57%	3.43%	0.58%	1.48%	1.51%
15	2.49%	3.75%	3.61%	2.25%	3.46%	3.34%	0.61%	1.49%	1.52%
16	2.44%	3.67%	3.54%	2.21%	3.40%	3.29%	0.61%	1.47%	1.51%
17	2.40%	3.60%	3.49%	2.16%	3.32%	3.21%	0.60%	1.44%	1.48%
18	2.37%	3.55%	3.44%	2.13%	3.26%	3.17%	0.59%	1.43%	1.44%
19	2.34%	3.50%	3.41%	2.09%	3.19%	3.13%	0.60%	1.42%	1.43%
20	2.31%	3.45%	3.38%	2.06%	3.16%	3.10%	0.57%	1.37%	1.39%

used within the optimization framework. In order to ensure that the results are robust, we suggest the following procedure. Take the solutions obtained by **EO-P*** and those obtained by **BM**, and compare them using *different* predictive models. Since different predictive models provide estimation using varied structure, it is not clear whether a solution that is superior for one predictive model will also be superior for other predictive models, built using alternative structural assumptions. If a solution is better across multiple predictive models, this makes the solution even more desirable.

Consider Tables 7 and 8. For both regions and each s , we use the solution from **EO-P*** and evaluate the number of sales predicted by the other predictive models, to see if the expansion decision remains favorable over **BM** when evaluated on another predictive model. This will enable managers to have more confidence in the solutions.

Let **LR** be the linear regression model, **LK** be the linear-kernel SVR model, and **RK** be the radial-kernel SVR model, where we include Wg as a predictor in the models only for Region 2. Each entry in the tables corresponds to the percent increase in additional expected sales realized from using the solution obtained by **EO-P*** over the solution obtained by **BM**. Interestingly, the solution found by **EO-P*** are almost always better than **BM** (i.e., most of the numbers in

Table 8: Robustness check for solution obtained by P^* for Region 2

K	$s = 2$			$s = 4$			$s = 6$		
	LR	LK	RK	LR	LK	RK	LR	LK	RK
1	6.07%	-2.50%	3.18%	5.69%	-4.68%	3.45%	4.21%	-5.40%	2.11%
2	7.64%	0.47%	4.43%	7.32%	-0.50%	4.21%	6.55%	-4.91%	2.26%
3	10.85%	-0.57%	5.50%	10.69%	-1.37%	4.54%	7.82%	-4.33%	2.43%
4	10.97%	-0.22%	5.63%	11.26%	-0.38%	5.64%	9.24%	-3.42%	2.67%
5	10.91%	-0.28%	5.51%	11.38%	-0.59%	5.54%	9.13%	-2.62%	2.54%
6	11.48%	0.39%	5.19%	12.38%	-0.12%	5.57%	8.57%	-2.12%	2.64%
7	10.65%	0.05%	4.88%	12.29%	-0.30%	5.98%	8.84%	-1.85%	2.58%
8	10.28%	0.22%	4.69%	11.80%	0.00%	5.67%	8.27%	-1.78%	2.43%
9	9.75%	0.43%	4.55%	11.59%	0.09%	5.46%	8.46%	-1.67%	2.62%
10	9.92%	0.45%	4.44%	11.80%	0.02%	5.32%	8.63%	-1.60%	2.78%
11	10.31%	0.36%	4.38%	11.95%	-0.11%	5.10%	8.31%	-1.09%	3.03%
12	13.53%	0.32%	4.50%	14.18%	-0.12%	5.08%	8.00%	-0.97%	2.96%
13	13.88%	0.32%	4.68%	14.50%	-0.03%	4.96%	7.77%	-0.97%	2.84%
14	13.35%	0.51%	4.53%	14.27%	0.13%	4.79%	8.39%	-0.86%	2.86%
15	12.82%	0.46%	4.39%	14.15%	0.19%	4.60%	8.27%	-0.81%	2.96%
16	12.60%	0.78%	4.28%	14.06%	0.40%	4.45%	8.00%	-0.68%	2.95%
17	13.16%	0.69%	4.27%	13.94%	0.50%	4.35%	7.89%	-0.57%	3.02%
18	13.34%	0.70%	4.22%	14.16%	0.61%	4.43%	7.50%	-0.54%	2.91%
19	13.01%	0.77%	4.33%	14.14%	0.73%	4.54%	7.36%	-0.48%	2.95%
20	13.96%	0.82%	4.62%	14.25%	0.70%	4.64%	7.22%	-0.47%	2.94%

Tables 7 and 8 are positive), independent of the predictive model that it is evaluated on. The exception is when we plug in **EO-P*** into **LK** in Region 2.

7 Conclusion and Future Work

In this paper, we provide a novel framework for jointly applying predictive modeling and optimization to expansion decisions for add-on products. Through collaboration with an industrial partner, we use point of sales data to generate predictive models and formulate optimization models over those predictive models to automate expansion decisions. Our results indicate that making expansion decisions using the framework described in this paper can result in substantial increases in expected sales over a baseline algorithm that only takes into account base-product sales.

This work has myriad possible directions for expansion. First, if given more complete and complex data, one could build more elaborate predictive models that provide a more fine-grained picture of what drives sales for add-on products. For example, a more detailed customer-level data on purchase behavior at different locations can be incorporated to develop improved

predictive models. This can also drive the prescriptive model, for instance, expanding to a new site frequented by existing customers may increase convenience for them but may limit market expansion to access new customers. Considering such nuances in customer behavior can lead to better location choice models. Second, these results should be replicated to more settings where expansion is limited by geographic constraints. Third, the dedicated solution methodology for solving optimization models resulting from non-linear predictive models should be investigated, so that expansion decisions are made optimally as opposed to by heuristics. Although the heuristic performs well for the company that we partnered with, there is no proof of optimality, and it is possible, and probable, that better collections of expansion sites exist.

References

- Aboolian R, Berman O, Krass D (2007a) Competitive facility location and design problem. *European Journal of Operational Research* 182(1):40 – 62, ISSN 0377-2217, URL <http://dx.doi.org/http://dx.doi.org/10.1016/j.ejor.2006.07.021>.
- Aboolian R, Berman O, Krass D (2007b) Competitive facility location model with concave demand. *European Journal of Operational Research* 181(2):598–619, URL <http://EconPapers.repec.org/RePEc:eee:ejores:v:181:y:2007:i:2:p:598-619>.
- Aboolian R, Berman O, Krass D (2009) Efficient solution approaches for a discrete multi-facility competitive interaction model. *Annals of Operations Research* 167(1):297–306, ISSN 1572-9338, URL <http://dx.doi.org/10.1007/s10479-008-0337-y>.
- ArcGIS (2017) How spatial autocorrelation (Global Moran's I) works. <http://pro.arcgis.com/en/pro-app/tool-reference/spatial-statistics/h-how-spatial-autocorrelation-moran-s-i-spatial-st.htm>, accessed: 2018-03-29.
- Aydin G, Ziya S (2008) Pricing promotional products under upselling. *Manufacturing & Service Operations Management* 10(3):360–376.
- Berman O, Drezner T, Drezner Z, Krass D (2009) Modeling competitive facility location problems: New approaches and results. *Decision Technologies and Applications*, 156–181 (INFORMS).
- Bertini M, Ofek E, Ariely D (2008) The impact of add-on features on consumer product evaluations. *Journal of Consumer Research* 36(1):17–28.
- Bertsimas D, Shioda R (2009) Algorithm for cardinality-constrained quadratic optimization. *Computational Optimization and Applications* 43(1):1–22.
- Bivand RS, Pebesma E, Gomes-Rubio V (2008) Applied spatial data analysis with R.
- Bozkaya B, Yanik S, Balcisoy S (2010) A gis-based optimization framework for competitive multi-facility location-routing problem. *Networks and Spatial Economics* 10(3):297–320, ISSN 1572-9427, URL <http://dx.doi.org/10.1007/s11067-009-9127-6>.
- Brown LA, Brown MA, Craig CS (1981) Innovation diffusion and entrepreneurial activity in a spatial context: Conceptual models and related case studies. *Research in Marketing: A Research Annual* 4:69–115.
- Church R, Murray A (2009) *Business site selection, location analysis, and GIS* (John Wiley & Sons), ISBN 9780470191064, URL <https://books.google.com/books?id=LgNNAQAAIAAJ>.

- Craig CS (1984) Models of the retail location process: a review. *Journal of Retailing* 60(1):5–36.
- Cui R, Gallino S, Moreno A, Zhang DJ (2017) The operational value of social media information. *Production and Operations Management* .
- Daskin M (2013) *Network and Discrete Location: Models, Algorithms, and Applications* (Wiley), ISBN 9781118536964, URL <https://books.google.es/books?id=L9I5n0Sym4cC>.
- Daskin MS (2008) What you should know about location modeling. *Naval Research Logistics (NRL)* 55(4):283–294, ISSN 1520-6750, URL <http://dx.doi.org/10.1002/nav.20284>.
- Dos Santos GF, Faria WR, et al. (2014) Spatial panel data models and fueldemand in brazil. *Anais do XL Encontro Nacional de Economia [Proceedings of the 40th Brazilian Economics Meeting]*, number 192 (ANPEC-Associação Nacional dos Centros de Pósgraduação em Economia [Brazilian Association of Graduate Programs in Economics]).
- Drezner T, Drezner Z, Kalczynski P (2015) A leaderfollower model for discrete competitive facility location. *Computers & Operations Research* 64:51–59, ISSN 0305-0548, URL <http://dx.doi.org/http://dx.doi.org/10.1016/j.cor.2015.04.012>.
- Drezner Z, Hamacher HW (2002) *Facility Location: Applications and Theory* (New York: Springer).
- Drezner Z, Zemel E (1992) Competitive location in the plane. *Annals of Operations Research* 40(1):173–193, ISSN 1572-9338, URL <http://dx.doi.org/10.1007/BF02060476>.
- Eiselt HA, Laporte G, Thisse JF (1993) Competitive location models: A framework and bibliography. *Transportation Science* 27(1):44–54, URL <http://dx.doi.org/10.1287/trsc.27.1.44>.
- Erat S, Bhaskaran SR (2012) Consumer mental accounts and implications to selling base products and add-ons. *Marketing Science* 31(5):801–818.
- Eze FJ, Odigbo BE, Ufot JA (2015) The correlation between business location and consumers patronage: Implications for business policy decisions. *British Journal of Economics, Management & Trade* 8(4):294–304.
- Ferreira KJ, Lee BHA, Simchi-Levi D (2015) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.
- Fisher M, Glaeser CK, Su X (2016) Optimal retail location: Empirical methodology and application to practice. *Available at SSRN*: URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2842064.
- Francis RL, White JA (1974) *Facility layout and location : an analytical approach*. Prentice-Hall

- international series in industrial and systems engineering (Englewood Cliffs, N.J. Prentice-Hall), ISBN 0-13-299149-7, URL <http://opac.inria.fr/record=b1083210>.
- Gerstner E, Hess JD (1990) Can bait and switch benefit consumers? *Marketing Science* 9(2):114–124.
- Getis A (2008) A history of the concept of spatial autocorrelation: A geographer’s perspective. *Geographical Analysis* 40(3):297–309.
- Ghosh A, Craig CS (1986) An approach to determining optimal locations for new services. *Journal of Marketing Research* 354–362.
- Hamacher HW, Nickel S, Schneider A (1996) Classification of location problems. Technical report, Location Science.
- Handler G, Mirchandani P (1979) *Theory of the Location of Industries* (Cambridge, MA: MIT Press).
- Hess JD, Gerstner E (1998) Yes, bait and switch really benefits consumers. *Marketing Science* 17(3):283–289.
- Hotelling H (1929) Stability in competition. *The Economic Journal* 39:41–57.
- Jain AK, Mahajan V (1979) Evaluating the competitive environment in retailing using multiplicative competitive interactive models. *Research in marketing* 2:217–235.
- Karakitsiou A (2015) *Modeling Discrete Competitive Facility Location* (Springer Publishing Company, Incorporated), 1st edition, ISBN 3319213407, 9783319213408.
- Ketelaar PE, Bernritter SF, van’t Riet J, Hühn AE, van Woudenberg TJ, Müller BC, Janssen L (2017) Disentangling location-based advertising: the effects of location congruency and medium type on consumers’ ad attention and brand choice. *International Journal of Advertising* 36(2):356–367.
- Kress D, Pesch E (2012) Sequential competitive location on networks. *European Journal of Operational Research* 217(3):483 – 499, ISSN 0377-2217, URL <http://dx.doi.org/http://dx.doi.org/10.1016/j.ejor.2011.06.036>.
- Kumar V, Leone RP (1988) Measuring the effect of retail store promotions on brand and store substitution. *Journal of Marketing Research* 178–185.
- Lin S (2017) Add-on policies under vertical differentiation: Why do luxury hotels charge for internet while economy hotels do not? *Marketing Science* .
- Love R, Morris J, Wesolowsky G (1988) *Facilities Location: Models & Methods*. Publications

- in Operations Research (North-Holland), ISBN 9780444010315, URL <https://books.google.com/books?id=FRpPAAAAAAAJ>.
- Moran PA (1950) Notes on continuous stochastic phenomena. *Biometrika* 37(1/2):17–23.
- Nishida M (2014) Estimating a model of strategic network choice: The convenience-store industry in okinawa. *Marketing Science* 34(1):20–38.
- Obeng E, Luchs R, Inman JJ, Hulland J (2016) Survival of the fittest: How competitive service overlap and retail format impact incumbents vulnerability to new entrants. *Journal of Retailing* 92(4):383–396.
- Öner Ö (2016) Retail productivity: The effects of market size and regional hierarchy. *Papers in Regional Science* .
- Pancras J, Sriram S, Kumar V (2012) Empirical investigation of retail expansion and cannibalization in a dynamic environment. *Management Science* 58(11):2001–2018.
- Rosenwein MB (1994) Discrete location theory, edited by p. b. mirchandani and r. l. francis, john wiley & sons, new york, 1990, 555 pp. *Networks* 24(2):124–125, ISSN 1097-0037, URL <http://dx.doi.org/10.1002/net.3230240212>.
- Verter V, Murat A (2007) S. nickel and j. puerto: Location theory: a unified approach. *Mathematical Methods of Operations Research* 66(2):369–371, URL <http://EconPapers.repec.org/RePEc:spr:mathme:v:66:y:2007:i:2:p:369-371>.
- Weber A, Über den Standort der Industrien, Tübingen (1909) *Theory of the Location of Industries* (Chicago: University of Chicago Press).
- Wilkie WL, Mela CF, Gundlach GT (1998a) Does bait and switch really benefit consumers? *Marketing Science* 17(3):273–282.
- Wilkie WL, Mela CF, Gundlach GT (1998b) Does bait and switch really benefit consumers? advancing the discussion. *Marketing Science* 17(3):290–293.
- Wolf GW (2011) Facility location: concepts, models, algorithms and case studies. series: Contributions to management science. *International Journal of Geographical Information Science* 25(2):331–333, URL <http://dblp.uni-trier.de/db/journals/gis/gis25.html#Wolf11>.