A Framework to Optimize Shelf Space Allocations for Physical Stores

Zhen Li
zhenmzn@amazon.com
Amazon
Seattle, WA, USA

Yuchen Luo
luoyuche@amazon.com
Amazon
Seattle, WA, USA

Gianluca Grilli
grillig@amazon.com
Amazon
Seattle, WA, USA

Xiaoyi Gu
guxiaoyi@amazon.com
Amazon
Seattle, WA, USA

Bruno Lopez-Videla
bllopezv@amazon.com
Amazon
Seattle, WA, USA

James Renier Domingo
rjamesre@amazon.com
Amazon
Seattle, WA, USA

ABSTRACT

In this paper we propose a novel optimization framework for store shelf space planning, specifically tailored for physical stores. The framework leverages machine learning and data-driven techniques to optimize shelf space allocation, aiming to maximize both short-term and long-term business metrics such as sales and profit. Our approach consists of three key components: a geospatial space elasticity model, a long-term causal lift model, and a space optimization module that incorporates various business constraints. To provide tailored recommendations for different store locations, we extend the classical space-dependent demand model by incorporating geospatial interventions and studying the geospatial heterogeneity effect on space elasticity, while considering inter-product substitutability and complementarity. Furthermore, our framework estimates the long-term impact of product assortment on store performance using customer transaction data. In case study experiments, our methodology showcased potential sales increases of 7.5% for store remodeling and 6.7% for new store openings compared to baseline approaches.

CCS CONCEPTS

• **Applied computing → Business intelligence**; **Decision analysis**; • **Computing methodologies → Machine learning approaches**.

KEYWORDS

retail optimization, shelf space planning, geospatial analysis, space elasticity, machine learning, causal inference, physical store operations, demand modeling, store layout optimization

ACM Reference Format:
Zhen Li, Yuchen Luo, Gianluca Grilli, Xiaoyi Gu, Bruno Lopez-Videla, and James Renier Domingo. 2025. A Framework to Optimize Shelf Space Allocations for Physical Stores. In *The 4th ACM SIGSPATIAL International Workshop on Spatial Big Data and AI for Industrial Applications (GeoIndustry '25), November 3–6, 2025, Minneapolis, MN, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3764919.3770880

GeoIndustry '25, November 3–6, 2025, Minneapolis, MN, USA

ACM ISBN 979-8-4007-2182-3/2025/11.
https://doi.org/10.1145/3764919.3770880

1 INTRODUCTION

Retail stores use hierarchical planning to optimize customer experience, with shelf space planning being crucial for store layout and product placement. This involves two main decisions: (1) how much space to allocate to each product family (e.g., department, category, or specific product) and (2) where to position these products within the store. Considerations include fixture allocation, shelf positioning, and strategic product placement, all of which affect customer demand and store performance.

This paper examines how to optimize shelf space allocation in retail stores. While traditionally based on manual decisions, modern retail now uses data-driven methods. Research introduced "space elasticity" - measuring how sales change relative to allocated space [4], which evolved from considering individual products [3] to including product relationships and display orientations [2, 6]. However, current methods overlook how location affects demand and space elasticity - urban stores show higher demand than suburban ones, with locations requiring different product focus (e.g., prepared foods in business districts, organic options in affluent areas).

Beyond location considerations, optimal space allocation must balance various costs (purchasing, inventory, salvage, shortage) with practical constraints like store footage limits and product-specific requirements. For example, Meat departments need careful space management to minimize waste, while alcohol sections face varying legal restrictions.

This paper proposes a framework for optimal retail shelf space allocation with three key contributions: 1) A space-dependent demand model incorporating geospatial variables for more accurate space elasticity estimates, especially for new stores, using a variant of Hierarchical Group Lasso [10] called Debiased Hierarchical Group Lasso; 2) A causal lift model measuring products' long-term impact on overall store sales (e.g., coffee's role in customer attraction); and 3) An optimization module combining space elasticity, long-term impact, costs, and business constraints. In simulation case studies conducted at Whole Foods Markets, the proposed framework demonstrated potential sales improvements of 7.5% for store remodeling processes and 6.7% for new store design processes compared to baseline methods.

2 SETUP AND PROBLEM FORMULATION

In the proposed framework, we consider space allocation of M product families within a store. Let $\boldsymbol{s} = (s_1, s_2, \ldots, s_M)$ denote allocated space, where s_i represents space for family i, $i = 1, \ldots, M$. Weekly unit sales $d_i(\boldsymbol{s}, \boldsymbol{x})$ depend on both its space s_i and other families' space $(s_j, j \neq i)$, with $\boldsymbol{x}$ representing geospatial attributes such as trade area demographics, foot and car traffic, and nearby grocery landscape. With average unit price p_i and K cost components $c_{i,k}(\boldsymbol{s}, \boldsymbol{x})$ (such as purchasing, shrinkage, inventory, and replenishment), store utility y is:

For revenue: $y = \sum_{i=1}^{M} d_i(\boldsymbol{s}, \boldsymbol{x}) \cdot p_i$. For profit: $y = \sum_{i=1}^{M} [d_i(\boldsymbol{s}, \boldsymbol{x}) \cdot p_i - \sum_{k=1}^{K} c_{i,k}(\boldsymbol{s}, \boldsymbol{x})]$.

Our objective is to determine $\boldsymbol{s}^* = \arg\max_{\boldsymbol{s}} y$, subject to business constraints.

3 METHODOLOGY

Our framework integrates two core models to optimize retail space allocation. Specifically, we employ a geospatial space elasticity model to estimate space elasticity and product demand, coupled with a causal lift model to assess the long-term impact of product families on overall store performance.

3.1 Geospatial space elasticity

The classical approach to study space elasticity is to model the demand for a product family based on the space allocated to each product family. Specifically, the demand for each product family is modeled as follows [2]:

$$d_i(\boldsymbol{s}) = \alpha_i \cdot s_i^{\beta_i} \cdot \prod_{j \neq i} s_j^{\delta_{ij}}, i = 1, \ldots, M. \tag{1}$$

Here, β_i denotes the primary space elasticity, measuring the magnitude of demand increase with the increase of its allocated space. The demand is assumed to grow at a diminishing rate, i.e., $0 \leq \beta_i \leq 1$, as illustrated by scientific experiments [5, 8]. δ_{ij} denotes the secondary space elasticity, which measures interdependency between product families i and j, ranging from -1 to 1, with positive values indicating complementary relationships and negative values indicating substitutes. α_i represents the base demand of product family i when all product families are allocated a unit space (i.e., $s_j = 1, j = 1, \ldots, M$).

Our methodology extends this traditional demand model (1) to incorporate geospatial variables $\boldsymbol{x}$, allowing the demand of a product family to depend not only on allocated space but also on relevant geospatial factors. Furthermore, we propose that space elasticity itself can be influenced by geospatial variables such as trade area population, foot and car traffic, and nearby grocery landscape. To be specific, our geospatial demand model is formulated as:

$$d_i(\boldsymbol{s}, \boldsymbol{x}) = \alpha_i \cdot s_i^{\beta_i} \cdot \prod_{j \neq i} s_j^{\delta_{ij}} \cdot \prod_{\ell=1}^{L} s_i^{\gamma_\ell x_\ell} \cdot \prod_{j \neq i} \prod_{\ell=1}^{L} s_j^{\zeta_{ij\ell} x_\ell} \cdot C(\boldsymbol{\tau}, \boldsymbol{x}), i = 1, \ldots, M. \tag{2}$$

In this extended model, γ_ℓ and $\zeta_{ij\ell}$ measures the geospatial heterogenity effect on the space elasticity from the geospatial variable x_ℓ, $\ell = 1, \ldots, L$. The primary space elasticity for product family i becomes $\beta_i + \sum_{\ell=1}^{L} \gamma_\ell x_\ell$, which is the exponential term of s_i. Therefore, if the shelf space of the target product family s_i increases by $r\%$, the demand $d_i(\boldsymbol{s}, \boldsymbol{x})$ will increase by $\left[(1 + r\%)^{\beta_i + \sum_{\ell=1}^{L} \gamma_\ell x_\ell} - 1\right] \times 100\%$. The secondary space elasticity is $\delta_{ij} + \sum_{\ell=1}^{L} \zeta_{ij\ell} x_\ell, j \neq i, j = 1, \ldots, M$. If the shelf space of the other product family j increases by $r\%$, the demand for the target product family $d_i(\boldsymbol{s}, \boldsymbol{x})$ will increase by $\left[(1 + r\%)^{\sum_{\ell=1}^{L} \delta_{ij} + \zeta_{ij\ell} x_\ell} - 1\right] \times 100\%$. $C(\boldsymbol{\tau}, \boldsymbol{x})$ measures the constant geospatial effect on demand, where $\log C(\boldsymbol{\tau}, \boldsymbol{x}) = \sum_{\ell=1}^{L} \tau_\ell x_\ell$. If we take logarithm for both sides of (2), it yields a linear equation

$$\begin{aligned} \log d_i(\boldsymbol{s}, \boldsymbol{x}) &= \log \alpha_i + \beta_i \log s_i + \sum_{j \neq i} \delta_{ij} \log s_j + \sum_{\ell=1}^{L} \gamma_\ell x_\ell \log s_i \\ &+ \sum_{j \neq i} \sum_{\ell=1}^{L} \zeta_{ij\ell} x_\ell \log s_j + \sum_{\ell=1}^{L} \tau_\ell x_\ell. \end{aligned} \tag{3}$$

To estimate the model parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\zeta}, \boldsymbol{\tau})$, where $\boldsymbol{\alpha} = \{\alpha_i\}_{1 \leq i \leq M}$, $\boldsymbol{\beta} = \{\beta_i\}_{1 \leq i \leq M}$, $\boldsymbol{\delta} = \{\delta_{ij}\}_{1 \leq i \neq j \leq M}$, $\boldsymbol{\gamma} = \{\gamma_\ell\}_{1 \leq \ell \leq L}$, $\boldsymbol{\zeta} = \{\zeta_{ij\ell}\}_{1 \leq i \neq j \leq M, 1 \leq \ell \leq L}$ and $\boldsymbol{\tau} = \{\tau_\ell\}_{1 \leq \ell \leq L}$, we propose a two-step process using a variant of Hierarchical Group Lasso (HGL) [10] and we name it Debiased Hierarchical Group Lasso (DHGL). This approach allows for variable selection on the high-dimensional geospatial variables $\boldsymbol{x}$ and robust parameter estimation for space elasticity. In the first step, we use HGL to perform variable selection, mitigating predictor multicollinearity and model overfitting. HGL ensures strong hierarchy in estimating interaction effects: when the interaction effect between a geospatial variable x_ℓ and allocated space s_i is non-zero, both the main effect of geospatial variables $(\tau_\ell x_\ell)$ and the interaction effect with allocated space $(\gamma_\ell x_\ell \log s_i)$ are retained after variable selection, which results in interpretable interaction models. Specifically, if $\gamma_\ell \neq 0$, we have $\tau_\ell \neq 0$. After variable selection, we obtain the remaining variable set F_i and rewrite (3) to

$$\begin{aligned} \log d_i(\boldsymbol{s}, \boldsymbol{x}) &= \log \alpha_i + \beta_i \log s_i + \sum_{\log s_j \in F_i} \delta_{ij} \log s_j \\ &+ \sum_{x_\ell \log s_i \in F_i} \gamma_\ell x_\ell \log s_i + \sum_{x_\ell \log s_j \in F_i, j \neq i} \zeta_{ij\ell} x_\ell \log s_j + \sum_{x_\ell \in F_i} \tau_\ell x_\ell. \end{aligned} \tag{4}$$

In the second step, to address potential bias in the parameter estimates provided by common regularization methods for linear models [15], we apply multiple linear regression on the selected variables F_i. This approach, often referred to as de-biased LASSO, yields debiased estimates for space elasticity. Such de-biased methods have been shown to possess desirable theoretical properties for estimator convergence in linear regression models [14, 17]. Note that the model observations for each product family depend on the number of stores and the duration of store operations (e.g. seasonal promotions, assortment changes), capturing both cross-sectional and longitudinal effects on the space elasticity.

3.2 Causal lift estimates

We develop a causal lift model to measure product families' long-term impact on store performance. Let $h_{i,T}(\boldsymbol{x})$ represent expected customer spend over T months following a purchase from family i, given location factors $\boldsymbol{x}$. Using the potential outcomes framework, we estimate Average Treatment Effect (ATE) [13, 7]:

$$h_{i,T}(\boldsymbol{x}) = \mathbb{E}[Y_i(1) - Y_i(0) | \boldsymbol{X} = \boldsymbol{x}] \quad (5)$$

where $Y_i(1)$ and $Y_i(0)$ are outcomes with and without purchase from family i, estimated using established methods such as the inverse probability weighted estimator (IPWE), doubly robust learning and Propensity Score Matching (PSM) [16, 9, 12, 11].

This analysis identifies product families that drive long-term customer value beyond their current sales performance. The resulting causal lift estimates complement our geospatial elasticity model by capturing temporal effects that extend beyond immediate space-sales relationships, directly informing our optimization module's balance between short-term and long-term utility maximization. For conciseness, this short paper omits detailed empirical results from our causal lift analysis across departments and focuses instead on the core space optimization framework.

3.3 Optimization and constraints

The optimization module ingests the aforementioned models and business constraints to generate the optimal space allocation for each product family, maximizing the store utility and enhancing customer experience. This utility consists of both short-term and long-term components. Short-term utility is primarily driven by store profit, which depends on unit sales $d_i(\boldsymbol{s}, \boldsymbol{x})$, price p_i and various cost components $c_{i,k}(\boldsymbol{s}, \boldsymbol{x})$. Long-term utility, on the other hand, is captured by causal lift estimates $h_{i,T}(\boldsymbol{x})$, which quantify the downstream impact of customer transactions following a purchase. To be specific, our objective is to solve the following optimization problem, subject to specific constraints:

$$\max_{\boldsymbol{s}} \sum_{i=1}^{M} \left\{ d_i(\boldsymbol{s}, \boldsymbol{x}) \cdot \left[\omega p_i + (1-\omega) h_{i,T}(\boldsymbol{x}) \right] - \sum_{k=1}^{K} c_{i,k}(\boldsymbol{s}, \boldsymbol{x}) \right\}. \quad (6)$$

In the objective function, p_i represents the average product price within a given product family, weighted by product unit sales in the store. For new stores, it is a common practice to estimate prices based on existing stores in the same geographical area, as prices tend to be relatively consistent within a region. The cost estimates $c_{i,k}(\boldsymbol{s}, \boldsymbol{x})$ can be developed using DHGL, the same approach used to estimate space elasticity in Section 3.1. $\omega \in [0, 1]$ represents a weight factor to balance the short- and long-term utility in the objective function.

The optimization problem is subject to several constraints commonly considered in space planning. These include (i) the sum of all product family spaces does not exceed the total store space limit, i.e., $\sum_{i=1}^{M} s_i \leq S$; (ii) setting lower and upper bounds for each product family's allocated space based on store layout or expert knowledge, i.e., $l_i \leq s_i \leq u_i$, $i = 1, \ldots, M$; (iii) limiting different cost components to certain thresholds, i.e., $c_{i,k}(\boldsymbol{x}, \boldsymbol{s}) \leq C_{i,k}$, $i = 1, \ldots, m$, $k = 1, \ldots, K$; and (iv) incorporating specific merchandising rules that reflect the company's strategy and consumer behavior for different product families [1]. This framework allows for a comprehensive approach to space allocation, balancing immediate profitability with long-term strategic considerations while adhering to practical operational constraints.

4 EMPIRICAL RESULTS

We evaluate our methodology through a case study of Whole Foods Market (WFM) stores, analyzing space elasticity's geospatial effects, departmental interdependencies, and applications to store remodeling and design.

4.1 Geospatial space elasticity and interdependency between product families

In Section 3.1, we introduce a location-dependent demand model (2) that accounts for the geospatial heterogeneity effect on space elasticity and the interdependency between product families. As an illustrative case study, we analyzed how the geospatial factors influence space elasticity for various departments in WFM stores (e.g., Produce, Meat, Prepared Foods). To achieve this, we leveraged DHGL to identify key variables with significant interaction effects on departmental space and demand, including trade area size, store tenure, consumables sales percentage, and the average home value in the trade area. Figure 1 shows the projected departmental weekly average sales (WAS) under actual and counterfactual primary elasticities when we increase the geospatial and store attribute variables by 10% for the Dairy and Prepared Foods departments in a representative urban store. The actual elasticity estimate is based on observed geospatial variables rather than counterfactual ones. If the counterfactual elasticity for a variable is higher than the actual elasticity estimate, it indicates the variable has a positive effect on the space elasticity. The results reveal that geospatial variables can have varying levels of effect on space elasticity across different departments. For instance, both trade area size and consumables spending in the trade area have a positive effect on space elasticity for the Dairy department, but the latter has a larger influence. Additionally, for Prepared Foods department, trade area size, home value, and tenure all show positive effects, with similar magnitudes of impact. These observations suggest the need for tailored recommendations for optimal shelf space allocations across different store locations. For a department with high space elasticity given a store location, the approach recommends allocating relatively large space to this department to improve store performance.

Furthermore, we studied the demand-space interdependency between departments in WFM stores. Figure 2 shows the mean estimates of primary and secondary elasticity for WFM departments, where the horizontal axis represents the allocated space and the vertical axis represents the department sales. The values reflect the magnitude of how the allocated space influences the demand for each department. We found that the space allocated to a department always has the highest effect on its own demand compared to the space allocated to other departments. The majority of departments do not significantly impact each other through allocated space. For the remaining department pairs, they exhibit complementary relationships with positive elasticity estimates to varying degrees. For example, the space allocated to the Meat department has a positive effect on the demand for several other departments, including Dairy, Produce, and Seafood, but the space allocated to the Bakery department does not influence the demand of any other departments.

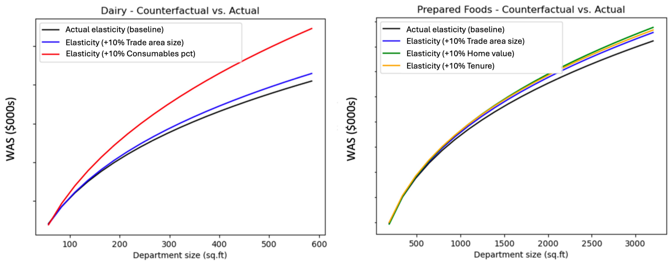

Figure 1: The projected weekly average sales (WAS) under actual and counterfactual primary elasticity estimates for department Dairy and Prepared Foods in a representative urban store. The elasticity value here is normalized to be the relative change of WAS when the department size increases by 10%. The WAS value in y-axis has been removed due to confidential purpose.

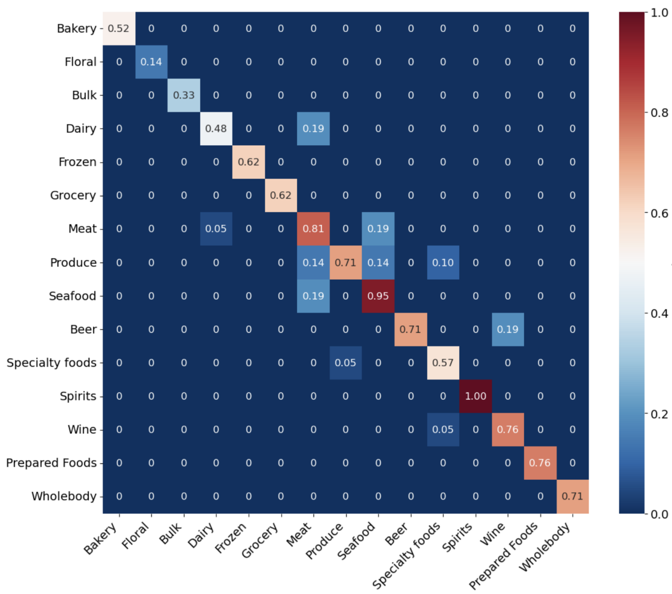

Figure 2: Mean estimates for primary and secondary space elasticity for WFM departments. X-axis represents the allocated space to departments and Y-axis represents the department sales. The values reflect the magnitude of how the allocated space influences demand for each department. All values have been normalized to a scale of 0 to 1 while preserving relative relationships.

4.2 Real-world Case Studies

Our framework was tested through case studies of optimizing space allocation for both existing store remodeling and new store designs at WFM, incorporating domain expertise as constraints.

The experiments incorporated realistic constraints derived from historical data. For remodeling scenarios, we applied reasonable bounds on department size variations. For new store scenarios, we implemented flexible constraints based on historical size distributions and cost parameters. Using these constraints, our experiments demonstrated potential sales improvements of 7.5% when tested across existing store remodeling scenarios and 6.7% when evaluated across new store design scenarios, compared to baseline approaches.

5 DISCUSSION AND CONCLUSION

This paper presents an optimization framework for retail shelf space allocation that maximizes store performance through a geospatial space elasticity model for location-specific recommendations, and a causal lift model that balances immediate profits with long-term impact. In our case studies, the proposed framework demonstrated the potential to achieve 7.5% and 6.7% higher sales for store remodeling and new store openings, respectively, compared to baseline methods.

Future research directions include: (1) optimizing product family locations within stores based on shopping behavior and complementarity (e.g., placing visually appealing items like Floral at entrances); (2) integrating omnichannel considerations, where online channels serve as virtual shelves and stores as fulfillment centers; and (3) coordinating shelf space with replenishment planning to balance space reduction with stock maintenance needs.

REFERENCES

[1] Teresa Bianchi-Aguiar, Maria Antónia Carravilla, and José F Oliveira. 2015. Replicating shelf space allocation solutions across retail stores. *Under revision on a peer reviewed journal.*

[2] Marcel Corstjens and Peter Doyle. 1983. A dynamic model for strategically allocating retail space. *Journal of the Operational Research Society*, 34, 10, 943–951.

[3] Marcel Corstjens and Peter Doyle. 1981. A model for optimizing retail space allocations. *Management Science*, 27, 7, 822–833.

[4] Ronald C Curhan. 1972. The relationship between shelf space and unit sales in supermarkets. *Journal of Marketing Research*, 9, 4, 406–412.

[5] Martin Eisend. 2014. Shelf space elasticity: a meta-analysis. *Journal of Retailing*, 90, 2, 168–181.

[6] Alexander Hübner and Kai Schaal. 2017. Effect of replenishment and backroom on retail shelf-space planning. *Business Research*, 10, 123–156.

[7] Guido W Imbens and Donald B Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences.* Cambridge University Press.

[8] Jens Irion, Jye-Chyi Lu, Faiz Al-Khayyal, and Yu-Chung Tsao. 2012. A piecewise linearization framework for retail shelf space management models. *European Journal of Operational Research*, 222, 1, 122–136.

[9] Zhen Li, Jie Chen, Eric Laber, Fang Liu, and Richard Baumgartner. 2023. Optimal treatment regimes: a review and empirical comparison. *International Statistical Review*, 91, 3, 427–463.

[10] Michael Lim and Trevor Hastie. 2015. Learning interactions via hierarchical group-lasso regularization. *Journal of Computational and Graphical Statistics*, 24, 3, 627–654.

[11] Paul R Rosenbaum and Donald B Rubin. 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39, 1, 33–38.

[12] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 1, 41–55.

[13] Donald B Rubin. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66, 5, 688.

[14] Sara A Van de Geer. 2008. High-dimensional generalized linear models and the lasso.

[15] Lu Xia, Bin Nan, and Yi Li. 2020. A revisit to de-biased lasso for generalized linear models. *arXiv preprint arXiv:2006.12778.*

[16] Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. 2012. A robust method for estimating optimal treatment regimes. *Biometrics*, 68, 4, 1010–1018.

[17] Cun-Hui Zhang and Stephanie S Zhang. 2014. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 76, 1, 217–242.