




Senseable City Lab :::: Massachusetts Institute of Technology

This paper might be a pre-copy-editing or a post-print author-produced .pdf of an article accepted for publication. For the definitive publisher-authenticated version, please refer directly to publishing house's archive system



Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Determinants of the localized behavior of individual food delivery choices[☆]

Rongxiang Su^{ID}*,¹, Javad Eshtiyagh^{ID}¹, Paolo Santi^{ID}, Songhua Hu^{ID},
Fábio Duarte^{ID}, Martina Mazzarello, Carlo Ratti

Senseable City Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

Dataset link: <https://github.com/rongxiangsu/talabat>

Keywords:

Online food delivery
Localized behavior
Localness
Nonlinear relationships
Explainable AI

ABSTRACT

Online food delivery systems have grown rapidly and globally due to urbanization, busy lifestyles, flexible work-life integration, and widespread mobile and Internet use. Despite their benefits, such express delivery services cause issues such as traffic congestion and increasing carbon emissions. The distance traveled to deliver food in the city is directly related to the transportation cost and environmental impact of delivery. Thus, containing this distance while satisfying citizens' needs for diverse food choices is central for urban sustainability. This study introduces a novel localness index that accounts for cuisine type to measure the degree of locality in food orders, which can serve as a proxy for the carbon efficiency of food delivery. Then, it uses a tree-based machine learning model and Explainable AI methods to offer insights into how dozens of variables influence the localness index of food delivery choices. Using Dubai as a case study, the analysis shows that current food delivery behavior is highly non-local and carbon inefficient. Such non-local behavior is mainly driven by variables such as delivery fee, cuisine type, total wait duration, the number and variety of available food options nearby, restaurant ratings, menu item prices, incentives, and socioeconomic conditions of the neighborhood. Furthermore, these factors may influence the locality of food orders in different, sometimes opposing ways across various areas of Dubai City, warranting specific attention from both the food delivery platform and local policymakers in order to encourage ordering from local restaurants and reduce emissions.

1. Introduction

The global online food delivery ecosystem has experienced remarkable growth, driven by increasing urbanization, flexible work-life integration, busy lifestyles, the proliferation of mobile devices, and widespread internet access, and has been recently boosted by the COVID-19 pandemic. This burgeoning industry offers consumers diverse culinary options from local favorites to international cuisines. However, this rapid expansion also presents challenges, including increased traffic congestion and emissions from delivery vehicles in urban areas (Li et al., 2020; Xie et al., 2021; Qin et al., 2021; Samudio Lezcano et al., 2023).

[☆] This article belongs to the Virtual Special Issue on “platform-based mobility”.

* Corresponding author.

E-mail address: rxsu@mit.edu (R. Su).

¹ These authors contributed equally to this work.

<https://doi.org/10.1016/j.trc.2025.105185>

Received 8 November 2024; Received in revised form 21 March 2025; Accepted 16 May 2025

Available online 25 June 2025

0968-090X/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

In certain regions, the rapid growth of online food ordering is more pronounced than in others. For example, in the Middle East and North Africa (MENA) region, the food delivery market was valued at approximately 9.8 million USD in 2022 and has been expanding at an annual rate of around 10%.² As a result, researchers and policymakers have been actively discussing various strategies to minimize costs and reduce environmental impact, including electrification (Löfstrand et al., 2013; Martins-Turner et al., 2020), alternative delivery modes such as autonomous vehicles (Figliozzi, 2020), and bundling strategies (Mancini and Gansterer, 2022). Beyond the exploration and implementation of these strategies, an equally important dimension is understanding individual food ordering behavior (Kimes, 2011; Yeo et al., 2017; Jun et al., 2021). Studying these behavioral patterns can provide insights into promoting more sustainable and efficient food delivery distribution and services.

Existing research primarily explores spatial, temporal, and behavioral patterns of online food delivery with a focus on demand (Kimes, 2011; Yeo et al., 2017; Jun et al., 2021; Wang and He, 2021) and the environmental impacts of online food delivery (Brown and Guiffreda, 2014; Van Loon et al., 2015; Xie et al., 2021; Long et al., 2023; Samudio Lezcano et al., 2023). Despite studies on the demand for and impact of food delivery, there appears to be a void in the literature in analyzing factors determining whether an individual orders from nearby or far away restaurants.

The spatial patterns of food delivery choices reflect the relation between objective and perceived food accessibility — as when customers order food from faraway restaurants even when similar options are available in the vicinity. Understanding factors affecting consumers' decisions regarding selecting local versus non-local options is key to developing more efficient public and business policies that reduce the total distance food travels in the city. Minimizing the total distance traveled will in turn lead to lower levels of carbon emissions and traffic. To capture this spatial behavior in food delivery, we propose a localness index that quantifies the extent to which an order is local with respect to the nearest equivalent offer at the time the order was placed.

Based on 2.7 million food delivery orders in Dubai, United Arab Emirates (UAE), in October 2023, this study investigates the determinants influencing the localness patterns of individual food delivery choices. Using a tree-based machine learning (ML) model and state-of-the-art Explainable AI methods (Lundberg et al., 2020), we examine the nonlinear effects of dozens of variables, such as the number and variety of food options available nearby, restaurant ratings, menu item prices, promotions or incentives each order received, and socioeconomic conditions of the neighborhood, on the localness of individual food delivery choice.

2. Literature review

Previous studies on food delivery have focused on two related aspects: the environmental impact of food delivery services and the factors affecting consumer choices on these platforms.

One of the primary environmental concerns related to food delivery services is carbon emissions. The increased use of delivery vehicles and motorcycles, many of which are powered by fossil fuels, significantly contributes to greenhouse gas emissions. While electric vehicles and motorcycles are becoming more common in food delivery services, they still contribute to environmental impacts indirectly due to the electricity consumption and the production of batteries required for their operation. According to Xie et al. (2021), each food delivery order in urban areas in China generates approximately 111.8 g of CO₂. Of this total, around 15.6 g is attributed to motorcycle-based food delivery, while 96.2 g is linked to food packaging. The results also demonstrate that by walking to a nearby restaurant to pick up instead of ordering food delivery, 68% of the CO₂ emissions derived from express food delivery could be reduced. Several other studies have analyzed the carbon emissions of food delivery in different contexts (Brown and Guiffreda, 2014; Van Loon et al., 2015; Li et al., 2020; Long et al., 2023). The obtained results consistently indicate that transporting food and last-mile delivery of express packages leads to significant levels of carbon emissions.

Consumer behavior on online food delivery platforms is a complex function of a large number of variables. Previous studies offer insight into some of these variables. Yeo et al. (2017) presents an overview of factors influencing individuals on online food delivery platforms by running surveys. The authors find the time saved, satisfaction with previous orders, and price savings to be the most influential factors. In another survey run by Jun et al. (2021), it is reported that the accuracy of detailed product information, detailed vendor information, perceived usefulness, and visual features and design play a significant role in persuading consumers to choose one restaurant over another. Following a similar methodology in the US, Kimes (2011) reports similar results but also highlights that location and type of cuisine are important factors that heavily influence consumers' choices.

In addition to survey data, a few pioneering studies leverage online food delivery usage information from large platforms. For instance, using online restaurant listings and monthly sales data collected from the major Chinese food delivery platform Ele.me, Wang and He (2021) explores the spatial patterns and determinants of online food delivery usage. Intuitively, the study finds that walk accessibility to food service and retail establishments, POI density, street intersection density, and the percentage of land used for rural residences, industries, and green spaces negatively impact online food delivery usage. Another study by Zhang et al. (2023), using data from Meituan, a Chinese food delivery company, discovers that while online food delivery generally improves the equity of food accessibility, it also exacerbates existing spatial disparities between urban and rural communities.

Several research gaps in the study of online food delivery warrant further investigation. First, most studies rely on small-scale survey data with limited information on the origin (restaurant location) and destination (customer location) of each delivery trip. While a few attempts have been made to acquire online food delivery data from local companies, no research to date has utilized detailed individual-level delivery data, including the origin and destination of each delivery trip. This information is crucial for

² <https://www.arabianbusiness.com/industries/travel-hospitality/revealed-uae-mena-witnessing-boom-in-restaurant-delivery-posts-10-jump-in-food-delivery-orders-in-q1-2024>.

understanding the heterogeneity in individual food ordering behavior. Second, to the best of our knowledge, existing studies have overlooked the local behavior of food delivery choices and the factors driving people to order locally. This is likely due to the limited access to granular information about the origin and destination of each delivery trip. Understanding local ordering behavior and its determinants has crucial implications for designing targeted policies, such as distributing incentives to specific neighborhoods or groups of customers and vendors to reduce delivery distances and emissions more effectively. This study addresses these gaps by developing a new index to measure the locality of food orders and examining the intricate relationships between various factors and the localness of food orders, using large-scale, detailed food delivery data with origin and destination information (i.e., longitude and latitude) provided by Talabat,³ a food delivery company based in Dubai, UAE.

3. Data

The primary data set in this study was obtained from Talabat, one of the largest food delivery companies in the Middle East. The data set spans from October 1 to October 31, 2023, capturing detailed information on customer orders over this period. It encompasses 2,778,722 orders which were delivered using motorcycles in Dubai. In total, this dataset contains information on 7,791 vendors and 555,860 unique users. Each order record includes key attributes such as the locations of the vendor and consumer, order placement timestamps, delivery timestamps, a list of purchased items including prices and descriptions, cuisine type, user IDs corresponding to multiple locations, basket amount, associated delivery fees, and discounts received.

The data set underwent several pre-processing steps to construct variables at the order, user, and neighborhood levels, aimed at analyzing factors influencing localized behavior. These steps include grouping similar cuisines into 66 distinct types (we discuss how the number of unique cuisine types may affect the analysis in the [Appendix](#)), calculating network distance from origins to destinations, neighborhood cuisine variety, order consistency, and estimating rental prices as a proxy of socioeconomic conditions. The rental price data were sourced from Bayut,⁴ one of the largest real estate portals in the UAE. The data provide information on approximately 32,000 rental units in Dubai. [Table 1](#) presents summary statistics of the input variables used in the model. Moreover, we also briefly describe the procedures for deriving some more complex variables in [Section 4.2](#). Summary statistics provide an initial overview of the data set: the average total wait duration for an order (measured from the time the order is placed until it is received) is 40.05 min, the average duration of delivery trips is 20.16 min, the average distance is 5.62 km, the average delivery fee is 4.8 dirhams, and the average order value is 66.51 dirhams.

4. Feature engineering

4.1. Prediction target: measure of localized behavior in individual food delivery choice

Building on recent work on the Local Mobility Index by [Malekzadeh et al. \(2024\)](#), we develop a localness index to measure the spatial behavior in individual food delivery choices. Originally, the Local Mobility Index was defined as the ratio of the distance from an anchor location to the chosen activity location to the distance from the anchor location to the nearest available activity of the same type ([Malekzadeh et al., 2024](#)). We adapt the Local Mobility Index to measure the extent to which an order is local relative to the nearest equivalent option at the time of ordering. The localness index is the ratio of the distance to the nearest open restaurant when the order is placed that offers the same type of cuisine as the chosen restaurant and the distance to the chosen restaurant, as formulated below. It is important to note that only nearby restaurants open at the same hour when the order is placed are considered in this process. In essence, it measures the level of localness of individuals' food delivery choices or the deviation from the nearest similar option relative to the customers' locations. The localness index L_i is bounded between 0 and 1, with values closer to 1 indicating more localized food delivery choices.

$$L_i = \frac{d_i^a}{d_i'^a} \quad (1)$$

where d_i^a is the network distance to the chosen restaurant of order i , $d_i'^a$ is the network distance to the nearest open restaurant offering the same cuisine type a from the customer's location of order i . Both d_i^a and $d_i'^a$ are incremented by extremely small values to prevent them from being zero. We implement network distance computation through the Pandana Python package ([Foti et al., 2012](#)) based on OpenStreetMap road network data.⁵

Although the localness index is derived from distance, it captures a crucial component that mere reference to distance cannot capture. For instance, if a person orders from a Chinese restaurant 4 km away, which is the nearest Chinese restaurant available, their order behavior is considered very local ($L = 1$). However, if another Chinese restaurant is only 1 km away and they still choose the one 4 km away, this order is considered less local ($L = 1/4$). Essentially, the localness index can differentiate between these behaviors effectively when the delivery distance is the same. The localness index can help identify areas where online food delivery orders are less localized, enabling urban planners to develop targeted strategies to encourage localized food ordering or suggest locations for new businesses.

³ <https://www.talabat.com/uae>.

⁴ <https://www.bayut.com/>.

⁵ <https://www.openstreetmap.org/>.

Table 1
Summary statistics of input variables (n = 2,778,722).

Variable	Description	Mean	Median	S.D.	Min	Max
<i>Order-level features</i>						
basket_amount	Total price (in dirham) for the order	66.51	52.77	51.98	0	3110
delivery_fee_amount	Delivery fee (in dirham) for the order	4.8	5.5	2.87	0	11
price_per_item	Average price per item (in dirham) from the chosen restaurant	28.3	28.11	14.64	0	201.11
order_time_hour	The hour of the day (0-23) when the order is placed	15.07	16	5.55	0	23
total_duration	Total wait duration of the order (in minutes) since it is placed	40.04	38	14.59	3	180
incentive_amount	Total incentives (in dirham) received for the order	7.69	0	13.27	0	1092
incentive_delivery	A dummy variable indicating if any delivery fee discount is applied (Yes = 1)	0.06				
incentive_menu_discount	A dummy variable indicating if any menu discount is applied (Yes = 1)	0.21				
incentive_voucher_wallet	A dummy variable indicating if any voucher wallet is applied (Yes = 1)	0.03				
isweekend	A dummy variable indicating if the order is on weekends (=1) or weekdays (=0)	0.31				
rating	Ordinal variable of the rating of the chosen restaurant (i.e., amazing, very good, good, OK; missing values are denoted as None)					
n_rating	Ordinal variable of the rating of the nearest open restaurant when the order is placed offering the same cuisine type					
main_cuisine	A categorical variable indicating cuisine types (66 unique cuisines in total)					
<i>User-level features</i>						
consistency	Individual ordering consistency measure (between 0 and 1) based on the entropy of the chosen cuisine types in the whole data for each customer	0.75	0.87	0.31	0	1
total_orders	Total number of orders placed by the same customer	13.78	10	13.69	1	239
customer_lat	Latitude of customers' locations					
customer_lon	Longitude of customers' locations					
<i>Neighborhood features</i>						
estimated_rent	Estimated neighborhood rent at each customer location	78.24	70.06	37.88	7.39	998.25
entropy_{x} km	Entropy (between 0 and 1) for measuring cuisine variety of open restaurants when the order is placed within an {x} km radius of customer location including 0.5, 1, 2, 3, 4, and 5 km.					
restaurant_count_{x} km	Total number of open restaurants when the order is placed within various radius intervals of customer location (within 0.5 km, between 0.5 and 1 km, 1 and 2 km, 2 and 3 km, 3 and 4 km, and 4 and 5 km)					

Notes: Empty cells indicate that the summary statistics are not applicable.

4.2. Input features: determinants of localness index

As illustrated in Fig. 1, our conceptual framework posits that individual localized food delivery ordering behavior is shaped by order-level, user-level, and neighborhood-level characteristics. Order-level features capture the most immediate aspects of a food delivery choice. User-level features capture individual differences and preferences in food ordering based on longitudinal observation of their historical orders. Neighborhood-level features represent the socioeconomic conditions around the customer's location, along with the availability and accessibility of food delivery services.

Specifically, for order-level features, we consider following factors: the total order price, delivery fee, the cuisine type of the selected restaurant, the average price per item from the selected restaurant, the hour the order is placed, total wait duration since the order is placed, incentive amounts, if the order receives certain discounts (e.g., delivery fee discount, menu discount, voucher wallet), whether the order is placed on a weekend day, the rating of the chosen restaurant, and the rating of the nearest open restaurant offering the same cuisine. At the user level, due to limited demographic information, we focus on the individual ordering consistency of cuisines based on one month of food delivery data, the total number of orders placed by the customer during the month, and the customer's location (i.e., longitude and latitude coordinates). For neighborhood-level features, we use neighborhood rent as a proxy for neighborhood socioeconomic conditions, alongside the variety of cuisines and the total number of open restaurants within several distance buffers around the customer's location when the order is placed.




Fig. 1. Conceptual framework for predicting the localized behavior of individual food delivery choices.

We control for feature multicollinearity by ensuring the Variance Inflation Factor (VIF) is below 4 for all features. Although machine learning models can handle multicollinearity when making predictions, this issue should not be neglected when interpreting models as many interpretation techniques, including the SHapley Additive exPlanations (SHAP) we employ in this study, assume independence among features (Hu et al., 2023). Eventually, 27 features (13 order-level, 4 user-level, and 10 neighborhood-level features) are computed and included in the final models. The descriptive statistics of each included predictor are summarized in Table 1. This section explains the procedures for deriving some more complex predictors.

(1) Individual consistency index of ordering behavior

We propose an individual consistency index based on entropy (Shannon, 1948) to measure the degree of consistency in the cuisine choices for each customer according to all their orders observed in the data. Originally, entropy was defined as a measure of uncertainty or disorder in a system (Shannon, 1948). Entropy is formulated as $H(X) = -\sum_{i=1}^n p(x_i) \log_b p(x_i)$. Here, entropy is adapted to quantify the consistency of individual food delivery choices. The individual consistency index is defined as

$$C_\alpha = \begin{cases} 0 & \text{if } n_\alpha = 1, \\ \frac{\sum_{j=1}^{n_\alpha} p_{\alpha,j} \ln(1/p_{\alpha,j})}{\ln(n_\alpha)} & \text{if } n_\alpha \neq 1. \end{cases} \quad (2)$$

where C_α is the consistency index based on the entropy of person α 's orders, $p_{\alpha,j}$ denotes a cuisine type j 's proportion of the total orders person α placed, and n_α is the total number of distinct cuisine types person α selected. The consistency index falls between 0 to 1. A value of 0 signifies no cuisine variety or high consistency, indicating the presence of only one cuisine type in person α 's orders. Conversely, a value of 1 indicates maximum cuisine variety, suggesting that all cuisines are equally represented, reflecting low consistency in cuisine choices.

(2) Rental prices as a proxy for socioeconomic status

Given the absence of direct socioeconomic status measurements in Dubai, we use home rental prices as a proxy for neighborhood socioeconomic status. Specifically, each customer location is assigned an estimated rental price based on the inverse distance-weighted average of the nearest twenty rentals, with an upper distance bound of 2 km. In the prediction model, we use price per square foot (dirhams/sqft) for unfurnished rental houses and apartments in the city.

(3) Dynamic entropy of neighborhood cuisine variety

Each customer location is assigned a dynamic entropy value representing the neighborhood's cuisine variety given the availability of open restaurants during the same hour when an order is placed. Building on the original formulation of entropy (Shannon, 1948), the dynamic entropy of neighborhood cuisine variety is defined as

$$E_{x(\alpha),t} = \frac{\sum_{j=1}^n p_{x(\alpha),t,j} \ln(1/p_{x(\alpha),t,j})}{\ln(n)} \quad (3)$$

where $E_{x(\alpha),t}$ is the entropy of neighborhood cuisine variety around a unique customer's location $x(\alpha)$ at time t (in hours from 0 to 23), $p_{x(\alpha),t,j}$ denotes a cuisine type j 's proportion of the total open restaurants within a buffer of customer's location $x(\alpha)$ at time t , and n is the fixed total number of distinct cuisine types in Dubai as observed in the data ($n = 66$). Several distance buffers around each customer's location are examined including 0.5, 1, 2, 3, 4, and 5 km. Notably, among these distance thresholds, 1, 3, and 4 km

have relatively high VIF ($VIF > 3$) so they are excluded from the final model. The dynamic entropy index of neighborhood cuisine variety ranges between 0 and 1. A value of 0 signifies no cuisine variety, indicating the presence of only one cuisine type around the customer's location $x(\alpha)$ at time t . Conversely, a value of 1 indicates maximum cuisine variety, suggesting that all cuisines are equally represented around the customer's location $x(\alpha)$ at time t .

(4) Number of nearby open restaurants

The number of open restaurants during the same hour when an order is placed within a spectrum of radius from each customer's location is used to measure the level of accessibility to restaurants at different times of the day. The distance buffers selected include 0.5, 1, 2, 3, 4, and 5 km. We use the difference in the total number of restaurants between two distance buffers to avoid multicollinearity. For instance, the restaurant count between 0.5 km and 1 km indicates the total number of open restaurants within this specific radius interval. As a result, the final model specification includes restaurant count within 0.5 km, restaurant count between 0.5 and 1 km, between 1 and 2 km, between 2 and 3 km, between 3 and 4 km, and between 4 and 5 km.

5. Methods: tree-based machine learning for nonlinear regression

Machine learning (ML) methods, particularly tree-based models such as Random Forests and Gradient Boosting Decision Trees (GBDT), have gained traction in various fields due to their ability to handle complex, nonlinear relationships that traditional statistical models struggle to capture (Bhavsar et al., 2017; Tizghadam et al., 2019; Tsolaki et al., 2023). Unlike conventional regression techniques, which assume a linear or predefined functional form between explanatory and dependent variables, ML models can automatically learn intricate patterns from data, making them highly effective in domains with heterogeneous and high-dimensional features (Boulesteix and Schmid, 2014). For example, in travel behavior research, tree-based ML models have been employed to analyze the nonlinear relationship between different classes of variables and mobility patterns (Ding et al., 2018; Tao et al., 2020; Yuan et al., 2022; Hu et al., 2023). These models accommodate threshold effects, where the marginal impact of an explanatory variable shifts significantly beyond a certain value (Galster, 2018). By leveraging machine learning's predictive power, researchers can uncover more nuanced insights that traditional models may overlook.

Despite capabilities in analyzing non-linear relationships, a common criticism of machine learning is its black-box nature, which has led to the growing emphasis on Explainable AI (XAI) methods. XAI techniques, such as SHapley Additive exPlanations (SHAP) and Partial Dependence Plots (PDP), enhance interpretability by quantifying the influence of individual variables on model outcomes (Holzinger et al., 2022; Samek et al., 2017; Molnar, 2022). In studies exploring food delivery behavior, for instance, ML models like CatBoost can effectively capture nonlinear effects of various order-level, user-level, and neighborhood-level factors, but XAI methods are crucial for making these insights actionable. Unlike traditional statistical analysis, which often provides global coefficients that assume uniform effects, XAI methods enable a more granular understanding of variable interactions, allowing researchers and policymakers to make data-driven decisions with greater confidence. By combining ML's flexibility with XAI's transparency, modern analytical approaches can bridge the gap between predictive accuracy and interpretability, offering a robust alternative to conventional statistical modeling.

5.1. CatBoost model

CatBoost (Prokhorenkova et al., 2018) is a member of the Gradient Boosting Decision Trees family (Friedman, 2001), renowned for its accuracy and robustness in tasks like ranking, classification, and regression. GBDT constructs an ensemble of decision trees sequentially, fitting each tree to the residual errors of the previous ones (Friedman, 2001). Unlike traditional GBDT, which struggles with high-cardinality categorical features, CatBoost efficiently handles them using an ordered target statistic strategy (Prokhorenkova et al., 2018). This capability is particularly valuable for this study, as it involves several categorical features, such as *main_cuisine*, which has 66 categories. Additionally, CatBoost's ordered boosting algorithm addresses prediction shifts caused by target leakage (Prokhorenkova et al., 2018). For technical details and implementation of CatBoost, the readers are referred to Prokhorenkova et al. (2018). The CatBoost model is implemented using *sklearn* (Pedregosa et al., 2011) and *catboost* (Prokhorenkova et al., 2018) Python libraries.

In the final CatBoost model, 2,778,722 delivery orders are included, among which 1,945,105 observations (70%) are used for training/validation with 5-fold cross-validation, and the remaining 833,617 observations (30%) are for testing. The six most important hyperparameters are tuned using *optuna* Python package (Akiba et al., 2019), including the number of trees (*iterations* = 1000), the minimum number of samples required to create a leaf (*min_data_in_leaf* = 11), the maximum depth of a tree (*depth* = 12), the learning rate (*learning_rate* = 0.285), the feature sampling rate (*subsample* = 0.736), and the percentage of features to use at each split selection (*colsample_bylevel* = 0.1). The Root Mean Square Error is used as the loss function. R^2 is used to evaluate model performance on the testing data set. Early stopping (*early_stopping_rounds* = 50) is used to prevent overfitting.

5.2. Interpretation approaches for ML models

The objective of interpretable ML is to elucidate how models generate predictions and to address questions such as the nature of the relationships between inputs and outputs and the identification of important features driving these predictions. In this study, we use two widely used model-agnostic interpretation techniques to understand the fine-tuned CatBoost model for predicting the localness of individual food delivery orders: the global interpretation method, Partial Dependence Plot (Friedman, 2001), and the local interpretation method, SHapley Additive exPlanations (Lundberg and Lee, 2017).

5.2.1. Global interpretation: Partial dependence plot

Partial Dependence Plot (PDP) is a post-hoc, model-agnostic interpretation technique that illustrates the average effect of changing one feature on the prediction across the data set (Friedman, 2001). PDP can be used to visualize and analyze the relationship between individual features and the response while marginalizing the values of all other features in the model. Formally, let $\mathbf{x} = \{x_1, x_2, \dots, x_p\}$ denotes the set of predictors in a model with the prediction function $\hat{f}(\mathbf{x})$. If $\mathbf{x}$ is divided into an interest set $\mathbf{z}_s$ and its complement $\mathbf{z}_c = \mathbf{x} \setminus \mathbf{z}_s$, the partial dependence of the response on $\mathbf{z}_s$ is defined as

$$f_s(\mathbf{z}_s) = \mathbb{E}_{\mathbf{z}_c} [\hat{f}(\mathbf{z}_s, \mathbf{z}_c)] = \int \hat{f}(\mathbf{z}_s, \mathbf{z}_c) p_c(\mathbf{z}_c) d\mathbf{z}_c \quad (4)$$

where $p_c(\mathbf{z}_c)$ is the marginal probability density of $\mathbf{z}_c$: $p_c(\mathbf{z}_c) = \int p(\mathbf{x}) d\mathbf{z}_s$. Due to its computational complexity, in practice, Eq. (4) is often approximated using training data set by

$$\hat{f}_s(\mathbf{z}_s) = \frac{1}{n} \sum_{i=1}^n \hat{f}(\mathbf{z}_s, \mathbf{z}_{i,c}) \quad (5)$$

where $\mathbf{z}_{i,c}$ ($i = 1, 2, \dots, n$) represents the values of $\mathbf{z}_c$ in the training set with total observations of n . Essentially, we average the effects of all the other predictors in the model.

5.2.2. Local interpretation: SHAP

One main limitation of PDP is their potential to obscure heterogeneous effects due to their computation of only average marginal effects across all observations. Additionally, outliers can distort the PD curve shape substantially, leading to misleading interpretations. An effective approach to overcome these issues is using local interpretation methods like SHAP (Lundberg and Lee, 2017). SHAP is derived from Shapley values in game theory, which was originally developed to quantify each player's contribution across all the permutations of the players solving a game (Shapley et al., 1953). In machine learning, the Shapley value of a feature measures its marginal contribution to the prediction compared to the average prediction for the data set, weighted and summed over all possible feature combinations. The Shapley value of a feature is defined as

$$\phi(X_k) = \sum_{S \subseteq \mathcal{X} \setminus X_k} \frac{|S|!(|\mathcal{X}| - |S| - 1)!}{|\mathcal{X}|!} (f(S \cup \{X_k\}) - f(S)) \quad (6)$$

$$f(S) = \int F(X) d\mathbb{P}_{\mathbf{x} \notin S} - E_X(F(X)) \quad (7)$$

where $\phi(X_k)$ is the Shapley value of feature X_k , $\mathcal{X}$ represents a set of all input features, S is the subset of $\mathcal{X}$ excluding X_k , $|\cdot|$ is the cardinality of a set, $f(S)$ is the prediction for features in subset S marginalized over features that are not included in S as formulated in Eq. (7), $f(S \cup \{X_k\})$ is the prediction for features in subset S plus feature X_k , $F(\cdot)$ is the approximation function of the estimator, and $\mathbb{P}_{\mathbf{x} \notin S}$ represents the distribution of features not in subset S .

SHAP explains the prediction of an instance by quantifying the contribution of each feature to the prediction based on Shapley values (Lundberg and Lee, 2017). In essence, each prediction is decomposed into several additive terms as shown below, each representing a certain feature's contribution to the model prediction.

$$\hat{y}_i = \phi_0 + \sum_{k=1}^{|\mathcal{X}|} \phi_i(X_k) \quad (8)$$

where $\hat{y}_i$ is the predicted value for the observation i , $\phi_0 = E(\hat{y})$ is the average prediction across all observations, and $\phi_i(X_k)$ is the SHAP value of feature X_k for the observation i .

The global SHAP importance of feature X_k is determined by calculating the absolute average of all individual Shapley values across data points.

$$I(X_k) = \frac{1}{n} \sum_{i=1}^n |\phi_i(X_k)| \quad (9)$$

where $I(X_k)$ is the importance of feature X_k and n is the total number of observations. Because of the high computational cost of SHAP, we use an efficient tree-specific variant of SHAP known as TreeExplainer (Lundberg et al., 2020) to interpret the fine-tuned CatBoost model.

SHAP enables both global and local explainability of the fine-tuned CatBoost model for predicting the localness of individual food orders. At the global level, collective SHAP values (Eq. (9)) help identify which predictors contribute the most to the overall prediction. At the local level, SHAP values for individual observations (Eq. (8)) reveal the heterogeneity in how different features impact the prediction of localness for specific orders.

5.2.3. Spatial interpretation of local nonlinear effects by SHAP

A common method for examining local effects is through SHAP dependence plots, which tend to be more informative for continuous features than categorical ones. However, these plots can often appear convoluted or random, failing to reveal clear patterns (see example results in our GitHub repository⁶). To enhance the interpretability of SHAP values, we develop a framework

⁶ https://github.com/rongxiangsu/talabat/blob/master/shap_plots/shap_beeswarm.pdf.




Fig. 2. Four quadrants of correlation between feature values and their SHAP values.

following Luo et al. (2025) that summarizes the correlation between feature values and their corresponding SHAP values, as shown in Fig. 2. Specifically, we divide these correlations into four quadrants: high feature values associated with positive SHAP (High-Positive), high feature values associated with negative SHAP (High-Negative), low feature values associated with positive SHAP (Low-Positive), and low feature values associated with negative SHAP (Low-Negative). High and low feature values can be defined using a threshold, such as the median or other meaningful values depending on the context of the study. We use the median as the threshold in subsequent analyses unless specified to ensure a more balanced data distribution across the four quadrants. For instance, the High-Positive category signifies that high feature values contribute to the prediction of localness. This classification framework enables us to map these four correlation types for spatial analysis.

6. Results

6.1. Localness patterns in food delivery choice

As Fig. 3(a) shows, the overall distribution of order travel distance is positively skewed. On average food delivery orders in Dubai travel 5.62 km to reach the customer's location. The distribution of the distance to the nearest open restaurant of the same cuisine type as the one chosen by the customer shows that, on average, a similar food option is located 1.66 km away from the customer's location. These results indicate that the average cost of "non-localness" is about 4 km. This means that, instead of ordering from the nearest restaurant at 1.66 km, people, on average, order from 5.62 km away — 3.4 times farther than necessary. Fig. 3(b) shows the results of the localness index for individual food orders. The localness index is generally low with a mean of 0.37 and positively skewed, with a notable concentration around the value of one. Only 5.06% of the orders have a localness index of one, meaning only these orders are placed from the nearest available restaurant. The results indicate that food delivery orders in Dubai are generally not local.

Regarding the time-of-day patterns of the localness index, Fig. 4 shows a gradual increase from midnight, reaching a peak around 5 to 6 am, followed by a decline until noon, after which it stabilizes. This pattern is likely due to the limited number of restaurants open after midnight, with the nearest option often being relatively far compared to daytime (see Fig. A.13 in the Appendix). As a result, people are more likely to order from or near the closest available options (i.e., higher localness) to receive their food more quickly. Fig. 5 illustrates the distribution of the localness index of the twenty least local cuisines. Food delivery orders for cuisines like Indian, fast food, BBQ, Arabic, and dessert tend to be less localized. Among these twenty least local cuisines, we observe that a smaller number of available restaurants for a given cuisine is likely correlated with a higher localness index. This is intuitive, as limited restaurant availability for a particular cuisine may limit people to order from local options. In contrast, Peruvian, Kuwaiti, Iraqi, Spanish, and Latin American cuisines exhibit the highest localness index of one. Among these high-localness cuisines, the number of available restaurants tends to be small (less than 10). Presumably, fewer restaurants offer these particular cuisines, leading people to order from the nearest available option.




Fig. 3. Distributions of (a) order travel distance (mean = 5.62 km, median = 4.72 km) and distance to the nearest open restaurant when the order is placed with the same cuisine as the chosen one (mean = 1.66 km, median = 1.03 km), and (b) localness of food delivery orders (mean = 0.37, median = 0.27). White triangles represent the mean.




Fig. 4. Time of day variation of localness index.

To further emphasize the significance of encouraging local food delivery, we conduct a simulated experiment by incrementally varying the localness of each order by 0.1 ($\delta_{localness}^+ = 0.1$) and recalculating the delivery distance manually according to Eq. (1). As shown in Fig. 6, increasing localness significantly reduces order travel distance, especially for long-distance orders. When localness reaches 1 (ordering from the nearest restaurant of the same cuisine), the majority of orders have a delivery distance below 4 km, with a mean of approximately 1.66 km (median 1.03 km). The average delivery distance decreases by 62.5%. Even a 0.1 increase in localness reduces the average distance from 5.62 km to 3.47 km. This suggests that increasing the localness of food delivery orders can significantly reduce food delivery distance, thereby decreasing vehicle kilometers traveled. By reducing the overall delivery distance, we can lower fuel consumption and decrease greenhouse gas emissions, contributing to a more sustainable delivery ecosystem.




Fig. 5. Distributions of localness index of twenty least local cuisines sorted by the mean. The number of available restaurants for each cuisine is noted in parentheses.




Fig. 6. Distributions of (a) estimated order travel distance (km) and (b) percentage reduction in order travel distance along with increasing localness of individual food delivery orders with 0.1 increments. The distribution at $\delta^+_{localness} = 0$ shows the original distribution of the food order travel distance.




Fig. 7. Global SHAP feature importance for predicting localness of individual food orders. Features are displayed on the vertical axis, sorted from most important at the top to least important at the bottom.

6.2. CatBoost regression results

As a result, the R^2 of the final fine-tuned CatBoost model is 0.74. This is a substantial improvement compared to the traditional Ordinary Least Squares (OLS) regression on the localness index with the same explanatory variables, which has an R^2 of only 0.299. The implementation and results of the OLS model are available in our GitHub repository.⁷

6.2.1. Feature importance

Fig. 7 illustrates the global SHAP feature importance for predicting the localness of food delivery orders. Among the tested features, *delivery_fee_amount* and *main_cuisine* rank as the top two important features boosting the model's predictions, with substantially higher SHAP values than all other features. The total duration of the order (*total_duration*), the number of open restaurants within a 500 m radius (*restaurant_count_0.5 km*) and between 500 m and 1 km (*restaurant_count_0.5_1 km*), the cuisine variety of restaurants within 500 m (*entropy_0.5 km*), ratings of the nearest open restaurant with the same cuisine type as the chosen restaurant (*n_rating*), and ratings of the chosen restaurant (*rating*) also play a key role in predicting the localness of food orders. The total SHAP values for thirteen order-level features, four user-level features, and ten neighborhood-level features are 0.317, 0.04, and 0.146, respectively. In essence, the selected order-level features play the most important role in predicting the localness of food delivery, followed by neighborhood features and user-level features.

6.2.2. Global nonlinear associations by PDP

In addition to examining feature importance, we also use Partial Dependence Plots to explore each feature's nonlinear effects on the prediction of localness. Fig. 8 illustrates the PDPs of the twelve order-level features in the fine-tuned CatBoost model. The graph of *main_cuisine* is not included in the paper due to its size but is available in our GitHub repository.⁸ It highlights the disparities in the association between different cuisines and the localness index.

⁷ https://github.com/rongxiangsu/talabat/blob/master/04_explore_localness.ipynb.

⁸ <https://github.com/rongxiangsu/talabat>.




Fig. 8. PDPs of the twelve order-level features in the fine-tuned CatBoost model. The y-axis represents the marginal effect of a feature on the predicted outcome.

First, delivery fee is a significant predictor for the localness index, exhibiting a positive correlation between 0 and 3.5 dirhams and a negative correlation beyond 3.5 dirhams. The positive correlation may be explained by the presence of a minimal delivery fee, even for orders from nearby restaurants. Additionally, customers with a Talabat membership subscription do not pay delivery fees, even when ordering from distant restaurants (low localness index). A negative correlation beyond 3.5 dirhams is expected, as higher delivery fees are typically associated with food orders from distant restaurants, indicating lower localness. Second, the total duration of the order (*total_duration*) has a clear declining log-linear relationship with localness, suggesting that localness decreases exponentially as the total duration increases. In essence, longer-duration orders tend to be less localized than shorter




Fig. 9. PDPs of two user-level features (*total_orders* and *consistency*) and one neighborhood-level feature (*estimated_rent*). Note that the PDP of the neighborhood-level feature *estimated_rent* is shown here for a concise presentation.

ones. Regarding the rating of the nearest same type of restaurant (*n_rating*), we find that when the rating is just “OK”, individual food delivery choices tend to be less localized, suggesting a preference for ordering from restaurants further away rather than the nearest option. Presumably, people avoid ordering from the nearest restaurants with “OK” ratings. The differences in the effects of other ratings on predicting localness are relatively minimal. Moreover, when the rating of the chosen restaurant (*rating* feature) is “amazing”, “very good”, or “OK”, the food delivery orders are more localized, indicating that individuals prefer nearby restaurants. Presumably, when ordering from nearby restaurants, people may either prioritize food quality by choosing highly rated options (i.e., “amazing” and “very good” ratings) or opt for lower-rated restaurants (i.e., “OK” ratings) to minimize wait time (e.g., fast food establishments often have relatively low ratings but offer quick preparation). Overall, the differences in the effects on localness among the three ratings—“amazing”, “very good”, and “OK”—are minimal.

The price per item of the customer-chosen restaurant (*price_per_item*) has a generally small positive effect on the localness index, meaning that orders from more expensive restaurants are relatively more localized. From the graph of *order_time_hour*, we observe that the overall localness of food delivery orders is substantially higher between 0 and 10 am compared to the rest of the day, which is consistent with the observation in Fig. 4. This is reasonable, as limited restaurants are open after midnight and the overall localness index between 0 and 10 am is higher than the rest of the day. Basket amount (*basket_amount*) has an almost monotonically positive effect on predicting localness, suggesting that expensive orders are more likely to come from nearby restaurants. However, the effect is small. As the graph of *isweekend* shows, food delivery orders are less localized on weekends compared to weekdays. However, the difference is small.

In terms of incentive-related features, we observe a clear threshold effect near 7 dirhams in terms of incentive amount (*incentive_amount*). Specifically, as customers receive incentives of up to 7 dirhams, they tend to order food less locally. This suggests that customers are more willing to order from further away when receiving incentives of up to 7 dirhams. However, the overall localness of food orders bounces back quickly when customers receive incentives of approximately 10 dirhams, and the effect of incentives on localness diminishes as customers receive more than 10 dirhams. This suggests that distributing more than 10 dirhams may have only a limited additional impact on encouraging local food orders. It is also noteworthy that the effect of incentive amount on localness is relatively small compared to other features such as *delivery_fee_amount* and *total_duration*. This is expected, given that only 37.3% of the orders receive incentives. When breaking down incentives by type—menu discounts (*incentive_menu_discount*), delivery discounts (*incentive_delivery*), and voucher wallets (*incentive_voucher_wallet*)—the first two tend to negatively impact the localness index, while the third enhances it. Menu-specific discounts generally have a negligible effect on localness. In contrast, delivery discounts significantly decrease the localness index, suggesting that when customers receive delivery fee discounts, they are more likely to order from distant restaurants rather than the nearest similar option. Meanwhile, voucher wallets positively influence local ordering behavior, indicating that when customers receive voucher wallet discounts, they are more likely to choose local restaurants.

Next, as shown in Fig. 9, the effects of the two user-level features, *total_orders* and *consistency*, are minimal. Overall, no clear patterns emerge in their influence on the localness index. The longitude and latitude of customers’ locations are included as predictors in the CatBoost model to capture spatial relationships for predicting the localness of food orders. However, the results of these two features are not reported here due to their limited interpretative value.

We next focus on the ten neighborhood-level features as illustrated in Figs. 9 and 10. As shown in the last graph in Fig. 9, the estimated rental price around customers’ locations (*estimated_rent*) has a non-linear effect on the model’s predictions, with substantial negative effects at lower values and diminishing effects as rent increases. This indicates that customers in less affluent neighborhoods tend to favor local orders more than their wealthier counterparts. In addition, as shown in Fig. 10, the impact of the dynamic entropy of cuisine variety within 500 m (*entropy_0.5_km*) and 2 km (*entropy_2_km*) on localness is generally negative, suggesting that customers from neighborhoods with more cuisine variety within 500 m or 2 km (higher entropy) tend to order from farther away restaurants than the nearest ones. One possible explanation is that these customers may prefer to pick up their orders, given the high variety of nearby restaurants within walking distance. However, since our data do not capture pickup orders, this behavior might




Fig. 10. PDPs of nine neighborhood features: *entropy_0.5 km*, *entropy_2 km*, *entropy_5 km*, *restaurant_count_0.5 km*, *restaurant_count_0.5_1 km*, *restaurant_count_1_2 km*, *restaurant_count_2_3 km*, *restaurant_count_3_4 km*, and *restaurant_count_4_5 km*.

be overlooked. Additionally, having a wide range of nearby dining options could lead to a sense of familiarity fatigue, prompting customers to seek non-local restaurants instead. Interestingly, we observe a positive effect of entropy within 5 km on the localness index, suggesting that greater cuisine variety within this range increases the likelihood of customers choosing local options.

We observe several important findings regarding the total number of open restaurants (regardless of cuisine types) within various radius intervals from customers' locations. First, the number of restaurants within 0.5 km, between 0.5 and 1 km, and between 1 and 2 km is overall negatively correlated with localness, suggesting that more restaurants near customers' locations within these distance ranges tend to discourage localized food orders. This is interesting and might even be a little counterintuitive. However, it is noteworthy that the food delivery data in this study only provide delivery orders, not in-person pick-up orders. Presumably, customers living in neighborhoods with more food supplies within walkable distance would pick up their food orders from the restaurants directly. For the rest of the features, there is a clear threshold effect: as the number of restaurants hits a certain value, the direction of the effect reverses. For instance, when the number of restaurants between 2 and 3 km is less than approximately 100, the overall impact is negative on localness. However, the impact becomes positive once the number exceeds 100 restaurants, suggesting that more restaurants between 2 and 3 km stimulate local ordering (although the positive effect is small). Overall, the effect size of the number of restaurants on localness decreases as the distance interval increases, which is intuitive.

6.2.3. Spatial interpretation of local nonlinear effects by SHAP

Previous discussions have primarily focused on the average effects of each selected feature in predicting the locality of food delivery orders through PDPs. In this section, we investigate the local nonlinear relationships between each feature and the predicted localness using SHAP. Given that the original dataset contains over 2 million data points, we aggregate the results by averaging both feature and SHAP values within hexagons of a size of approximately 0.737 km² (with an edge length of 0.531 km). We select several important continuous features for discussions, including estimated rent near the customer, cuisine variety within 0.5 km, 2

km, and 5 km of the customer's location, incentive amount, the proportion of orders receiving delivery discounts, and the number of restaurants within various distance ranges: 0.5 km, 0.5–1 km, 1–2 km, 2–3 km, 3–4 km, and 4–5 km from the customer. The SHAP maps for the rest of the continuous features are available in our GitHub repository.

As shown in Fig. 11, the first graph illustrates the correlation between estimated rent near the customer and its SHAP values. For this feature, we use the first and fourth quartiles of estimated rent to distinguish between low and high rents, highlighting several significant neighborhoods. Only a few areas appear in orange, representing neighborhoods with rents in the fourth quartile (higher rents) that are associated with low localness. In contrast, the brown areas indicate neighborhoods where higher rents correspond to more local ordering behavior. Some examples of High-Positive correlations are found in areas such as the Palm, Jumeirah area, and Dubai Marina. These patterns reveal the contrasting effects of neighborhood rent on local food delivery behavior. In low-rent areas, we observe a fairly balanced distribution of hexagons displaying either a Low-Negative (white) or Low-Positive (blue) association.

Regarding *entropy_0.5 km*, inner-city areas predominantly display High-Positive and High-Negative associations. Interestingly, these areas are surrounded by regions with Low-Positive associations. Note that there is no Low-Negative association between *entropy_0.5 km* and localness. Overall, the spatial patterns suggest that with greater cuisine variety within 500 m of customer locations, people in inner city areas tend to order either more or less locally. However, in low-variety neighborhoods, individual food delivery choices consistently remain more local. We observe a prevalence of Low-Positive patterns for the feature *entropy_2 km*, indicating that in these blue Low-Positive areas, people experience less cuisine variety within 2 km but tend to order more locally. In the graph of *entropy_5 km*, there is a significantly larger number of Low-Negative areas compared to the previous two features, indicating that these regions have low cuisine variety within 5 km of customer locations, which correlates with less localized ordering behavior.

In terms of total incentive amount, most areas exhibit a High-Negative association, suggesting that offering incentives in these regions discourages local food orders. Only a few brown hexagons show a High-Positive correlation. These results imply that if the goal is to encourage ordering from local restaurants, platforms like Talabat or local policymakers may need to reconsider their incentive strategies for the targeted High-Negative orange areas. Interestingly, most areas show a Low-Positive correlation when focusing on the proportion of orders receiving delivery discounts (*incentive_delivery*). Some areas exhibit a High-Positive correlation, suggesting that distributing incentives in these regions can effectively encourage more local ordering. In contrast, High-Negative areas are less effective, as distributing incentives there tends to reduce localness. The results for other incentive types, such as voucher wallets and menu discounts, are available in our GitHub repository.

The remaining graphs focus on the correlation between the number of restaurants within different distance intervals from customer locations and the localness of food delivery. Several interesting patterns emerge. First, only a few orange areas, representing higher restaurant access within 500 m, consistently show a High-Negative association with localness. This suggests that, despite having nearby options, people in these areas prefer ordering from distant restaurants. Alternatively, they may be more inclined to pick up orders due to the high accessibility of restaurants within walking distance. However, since our data do not record pickup orders, this behavior may not be observable. Second, blue areas indicate a Low-Positive association, meaning that food orders are more localized in areas with fewer nearby restaurants. These areas are primarily situated farther from city centers. Third, we observe a few hexagons with High-Positive correlations (brown color), suggesting that areas with greater restaurant access within a specific distance interval tend to have more localized orders. This phenomenon is particularly evident for the feature *restaurant_count_0.5 km* and *restaurant_count_0.5_1 km*, suggesting that people living in these brown hexagon areas likely have good access to food delivery options and tend to place orders from nearby restaurants. Note that there is no Low-Negative association for the features *restaurant_count_0.5 km* and *restaurant_count_0.5_1 km*, suggesting that people in areas with limited food delivery access within 500 m or between 500 m and 1 km are more likely to order from the nearest options.

7. Discussion and conclusions

This study proposes a localness index to measure the extent to which an order is local with respect to the nearest equivalent offer at the time the order was placed. Using food delivery order data in Dubai in October 2023 as a case study, we find that food delivery orders in Dubai are not local, with only 5.06% of the orders from the nearest available restaurants and a large concentration of orders being extremely far from the nearest similar restaurants. Current food delivery patterns in Dubai impose an environmental cost of more than 3.4 times higher than the local, carbon-efficient delivery pattern. Furthermore, using a tree-based machine learning model and two model-agnostic Explainable AI techniques (i.e., PDP and SHAP), this study identifies several important nonlinear effects of order-level and neighborhood-level features on predicting the localness of individual food delivery orders. The identified non-linear relationships offer a more nuanced perspective than the typical linear models commonly used in previous research. This distinction is crucial because linear models often oversimplify the complexity of factors influencing consumer behavior, failing to capture the variations observed in real-world data. We discuss the key findings and their implications below.

7.1. Key findings

First, there is an inverse relationship between delivery fee and localness beyond 3.5 dirhams, supporting the intuition that higher delivery fees correspond to more distant restaurant choices. Platforms like Talabat could introduce dynamic pricing that rewards local choices (e.g., discounts or loyalty points for ordering nearby). Delivery fee structures could be adjusted to reflect sustainability




Fig. 11. Spatial distribution of the correlation between feature values and their SHAP values, aggregated at the hexagon level.




Fig. 11. (continued).

goals—lower fees for short distances and progressively higher fees for long distances to discourage unnecessary long-haul orders.

Second, we observe a clear threshold effect of incentive amount on the prediction of localness near seven and ten dirhams. Specifically, as customers receive incentives of up to seven dirhams, they tend to order food less locally. However, the overall localness bounces back quickly when customers receive incentives of approximately ten dirhams, and the effect of incentives on localness diminishes as customers receive more than ten dirhams. This finding has significant implications for policymaking, particularly in the context of incentivizing consumer behavior in urban environments. The observed threshold effect suggests that policymakers and businesses could tailor incentives more effectively based on this nonlinear relationship. For example, offering small incentives (below seven dirhams) may not significantly encourage local ordering, as consumers are less likely to change their behavior in response to modest incentives. On the other hand, increasing the incentive to around ten dirhams could have a much more pronounced effect, promoting local food purchases and supporting local businesses. However, offering incentives greater than ten dirhams may not yield a proportional increase in localness, suggesting diminishing returns at higher levels. This insight could help food delivery platforms design incentive structures that maximize local economic benefits and reduce transportation emissions from food delivery. In addition, the voucher wallet appears to be the most effective incentive for promoting local ordering and sustainable food delivery compared to menu discounts or delivery discounts.

Third, we observe a negative effect of cuisine variety within 500 m and 2 km on localness suggesting that customers in areas with high nearby variety may be more inclined to order from farther restaurants. Policymakers could consider incentives for local food delivery, such as loyalty programs for ordering from nearby restaurants. This could encourage more sustainable food consumption patterns. If variety fatigue leads customers to order from farther restaurants, dynamic marketing strategies — such as rotating featured local restaurants or offering discounts for trying new nearby options — could encourage customers to rediscover local businesses instead of defaulting to distant choices.

Last but not least, policies aimed at reducing emissions might use targeted incentives in particular neighborhoods to nudge consumers toward local ordering, aligning environmental goals with consumer behavior in a more efficient manner. To achieve this, the spatial analysis of SHAP allows us to identify how factors like incentive amount, neighborhood rent, cuisine variety, and restaurant density impact local food delivery behavior across different areas in the city. Unlike most existing studies that lack spatially detailed data, our findings highlight specific regions in Dubai where tailored incentives could encourage more local ordering. This aligns with recent calls in the food delivery literature for more targeted strategies that account for spatial heterogeneity in consumer behavior (Wang and He, 2021; Zhang et al., 2023).

7.2. Limitations and future work

This work has several limitations. First, because the effects of various features on localness behavior are highly specific to local contexts in Dubai, additional studies are needed to explore localness behavior in other cities and under different socio-economic and cultural conditions to better understand its broader applicability and implications (Minaker et al., 2016). Second, like many existing studies, the present interpretation of tree-based ML models cannot establish causality between variables; it only explains the nonlinear associations between them. Future research may explore causal inference methods to determine if causation exists among key predictors identified, such as cuisine variety entropy, number of nearby restaurants, and incentives. Third, due to data unavailability, this study considers only one socioeconomic feature, the rental price. Future work will incorporate additional socioeconomic and demographic features, either at a neighborhood or individual level, once relevant data become available. Fourth, this study relies on only one month of food delivery data, limiting our ability to capture seasonal effects on the localness index and variations in usage patterns throughout the year. In addition, variations in the scale of food delivery services across restaurants may influence the localness index. In future work, we will take these factors into account. Lastly, the localness index is sensitive to how cuisine types are categorized (see the Appendix), which may influence the nonlinear relationships discussed. Future work could refine cuisine categorization and systematically assess the robustness of the findings.

CRediT authorship contribution statement

Rongxiang Su: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Javad Eshtiyagh:** Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Paolo Santi:** Writing – review & editing, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **Songhua Hu:** Writing – review & editing, Methodology, Investigation, Conceptualization. **Fábio Duarte:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Martina Mazzarello:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Carlo Ratti:** Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors gratefully acknowledge Talabat company (<https://www.talabat.com/uae>) for providing the food delivery transaction data and the members of the MIT Senseable City Lab consortium. Special thanks to Dr. Peng Luo (Postdoc fellow at MIT Senseable City Lab) for his valuable suggestions on improving the spatial interpretation of SHAP. The authors also appreciate the comments by the anonymous reviewers and the editor that improved this paper substantially.

Appendix

As Fig. A.12 shows, the box plot illustrates the sensitivity of the localness index to the number of cuisine type categories used in its calculation. The original dataset contains 123 unique cuisine types. In this analysis, we tested different cuisine categorizations for localness index calculation. The dictionary for each cuisine categorization is available on our Github repository.⁹ As the number of cuisine types increases from 10 to 90, the median and interquartile range of the localness index also increase, suggesting that a finer categorization captures greater variation in localness. When only 10 cuisine types are used, the localness index tends to be lower, likely due to the aggregation of diverse cuisines into broad categories, which may obscure distinctions between local and non-local establishments. Conversely, as the number of cuisine types increases the localness index becomes bigger. This analysis shows that different classification schemes of cuisine types can influence the localness index calculation.






Fig. A.12. Sensitivity analysis of localness index to different cuisine categorizations. White triangles represent the mean.

⁹ https://github.com/rongxiangsu/talabat/blob/master/different_cuisine_classifications.py.



(a)



(b)

Fig. A.13. Time of day variation of (a) order travel distance and (b) the distance to the nearest open restaurant offering the same cuisine type as the chosen restaurant. White triangles represent the mean.

Data availability

The main data for this study come from Talabat's food delivery transaction records. The authors do not have permission to share the data. Our code is accessible in a GitHub repository <https://github.com/rongxiangsu/talabat>.

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M., 2019. Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 2623–2631.
- Bhavsar, P., Safro, I., Bouaynaya, N., Polikar, R., Dera, D., 2017. Machine learning in transportation data analytics. In: Data Analytics for Intelligent Transportation Systems. Elsevier, pp. 283–307. <http://dx.doi.org/10.1016/B978-0-12-809715-1.00012-2>, URL: <https://linkinghub.elsevier.com/retrieve/pii/B9780128097151000122>.

- Boulesteix, A.-L., Schmid, M., 2014. Machine learning versus statistical modeling. *Biom. J.* 56 (4), 588–593. <http://dx.doi.org/10.1002/bimj.201300226>, URL: <https://onlinelibrary.wiley.com/doi/10.1002/bimj.201300226>.
- Brown, J.R., Guiffrida, A.L., 2014. Carbon emissions comparison of last mile delivery versus customer pickup. *Int. J. Logist. Res. Appl.* 17 (6), 503–521, Publisher: Taylor & Francis.
- Ding, C., Cao, X.J., Næss, P., 2018. Applying gradient boosting decision trees to examine non-linear effects of the built environment on driving distance in Oslo. *Transp. Res. Part A: Policy Pr.* 110, 107–117. <http://dx.doi.org/10.1016/j.tra.2018.02.009>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0965856417310030>.
- Figliozzi, M.A., 2020. Carbon emissions reductions in last mile and grocery deliveries utilizing air and ground autonomous vehicles. *Transp. Res. Part D: Transp. Environ.* 85, 102443. <http://dx.doi.org/10.1016/j.trd.2020.102443>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1361920920306301>.
- Foti, F., Waddell, P., Luxen, D., 2012. A generalized computational framework for accessibility: from the pedestrian to the metropolitan scale. In: *Proceedings of the 4th TRB Conference on Innovations in Travel Modeling*. Transportation Research Board. pp. 1–14.
- Friedman, J.H., 2001. Greedy function approximation: A gradient boosting machine. *Ann. Statist.* 29 (5), <http://dx.doi.org/10.1214/aos/1013203451>, URL: <https://projecteuclid.org/journals/annals-of-statistics/volume-29/issue-5/Greedy-function-approximation-A-gradient-boosting-machine/10.1214/aos/1013203451.full>.
- Galster, G.C., 2018. Nonlinear and threshold effects related to neighborhood: Implications for planning and policy. *J. Plan. Lit.* 33 (4), 492–508. <http://dx.doi.org/10.1177/0885412218793693>, URL: <http://journals.sagepub.com/doi/10.1177/0885412218793693>.
- Holzinger, A., Goebel, R., Fong, R., Moon, T., Müller, K.-R., Samek, W. (Eds.), 2022. xxAI - Beyond explainable AI: International workshop, held in conjunction with ICML 2020, July 18, 2020, Vienna, Austria, revised and extended papers. In: *Lecture Notes in Computer Science*, vol. 13200, Springer International Publishing, Cham, <http://dx.doi.org/10.1007/978-3-031-04083-2>, URL: <https://link.springer.com/10.1007/978-3-031-04083-2>.
- Hu, S., Xiong, C., Chen, P., Schonfeld, P., 2023. Examining nonlinearity in population inflow estimation using big data: An empirical comparison of explainable machine learning models. *Transp. Res. Part A: Policy Pr.* 174, 103743. <http://dx.doi.org/10.1016/j.tra.2023.103743>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0965856423001635>.
- Jun, K., Yoon, B., Lee, S., Lee, D.-S., 2021. Factors influencing customer decisions to use online food delivery service during the COVID-19 pandemic. *Foods* 11 (1), 64. <http://dx.doi.org/10.3390/foods11010064>, URL: <https://www.mdpi.com/2304-8158/11/1/64>.
- Kimes, S.E., 2011. Customer perceptions of electronic food ordering.
- Li, C., Mirosa, M., Bremer, P., 2020. Review of online food delivery platforms and their impacts on sustainability. *Sustain.* 12 (14), 5528. <http://dx.doi.org/10.3390/su12145528>, URL: <https://www.mdpi.com/2071-1050/12/14/5528>.
- Löfstrand, S., Hellgren, J., Thulin, N., Engdahl, H., Pettersson, S., Williamsson, J., Sandoff, A., 2013. Feasibility of electrifying urban goods distribution trucks. *SAE Int. J. Commer. Veh.* 6 (2013-01-0504), 24–33.
- Long, Y., Chen, G., Wang, Y., 2023. Carbon footprint of residents' online consumption in China. *Environ. Impact Assess. Rev.* 103, 107228. <http://dx.doi.org/10.1016/j.eiar.2023.107228>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0195925523001944>.
- Lundberg, S.M., Erion, G., Chen, H., DeGrave, A., Prutkin, J.M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., Lee, S.-I., 2020. From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* 2 (1), 56–67. <http://dx.doi.org/10.1038/s42256-019-0138-9>, URL: <https://www.nature.com/articles/s42256-019-0138-9>.
- Lundberg, S.M., Lee, S.-I., 2017. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* 30.
- Luo, P., Chen, C., Gao, S., Zhang, X., Majok Chol, D., Yang, Z., Meng, L., 2025. Understanding of the predictability and uncertainty in population distributions empowered by visual analytics. *Int. J. Geogr. Inf. Sci.* 39 (3), 675–705, Publisher: Taylor & Francis.
- Malekzadeh, M., Reuschke, D., Long, J.A., 2024. Quantifying local mobility patterns in urban human mobility data. *Int. J. Geogr. Inf. Sci.* 1–24. <http://dx.doi.org/10.1080/13658816.2024.2389410>, URL: <https://www.tandfonline.com/doi/full/10.1080/13658816.2024.2389410>.
- Mancini, S., Gansterer, M., 2022. Bundle generation for last-mile delivery with occasional drivers. *Omega* 108, 102582, Publisher: Elsevier.
- Martins-Turner, K., Grahle, A., Nagel, K., Göhlich, D., 2020. Electrification of urban freight transport-a case study of the food retailing industry. *Procedia Comput. Sci.* 170, 757–763, Publisher: Elsevier.
- Minaker, L.M., Shuh, A., Olstad, D.L., Engler-Stringer, R., Black, J.L., Mah, C.L., 2016. Retail food environments research in Canada: A scoping review. *Can. J. Public Heal.* 107 (Suppl 1), eS4–eS13, Publisher: Springer.
- Molnar, C., 2022. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*, second ed. URL: <https://christophm.github.io/interpretable-ml-book>.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A., 2018. CatBoost: unbiased boosting with categorical features. *Adv. Neural Inf. Process. Syst.* 31.
- Qin, H., Wei, Y., Zhang, Q., Ma, L., 2021. An observational study on the risk behaviors of electric bicycle riders performing meal delivery at urban intersections in China. *Transp. Res. Part F: Traffic Psychol. Behav.* 79, 107–117. <http://dx.doi.org/10.1016/j.trf.2021.04.010>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1369847821000917>.
- Samek, W., Wiegand, T., Müller, K.-R., 2017. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*.
- Samudio Lezcano, M., Harper, C.D., Nock, D., Lowry, G.V., Michalek, J.J., 2023. Online grocery delivery: Sustainable practice, or congestion generator and environmental burden? *Transp. Res. Part D: Transp. Environ.* 119, 103722. <http://dx.doi.org/10.1016/j.trd.2023.103722>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1361920923001190>.
- Shannon, C.E., 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27 (3), 379–423, Publisher: Nokia Bell Labs.
- Shapley, L.S., et al., 1953. *A Value for N-Person Games*. Publisher: Princeton University Press Princeton.
- Tao, T., Wang, J., Cao, X., 2020. Exploring the non-linear associations between spatial attributes and walking distance to transit. *J. Transp. Geogr.* 82, 102560. <http://dx.doi.org/10.1016/j.jtrangeo.2019.102560>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0966692318307233>.
- Tizghadam, A., Khazaei, H., Moghaddam, M.H.Y., Hassan, Y., 2019. Machine learning in transportation. *J. Adv. Transp.* 2019, 1–3. <http://dx.doi.org/10.1155/2019/4359785>, URL: <https://www.hindawi.com/journals/jat/2019/4359785/>.
- Tsolaki, K., Vafeiadis, T., Nizamis, A., Ioannidis, D., Tzovaras, D., 2023. Utilizing machine learning on freight transportation and logistics applications: A review. *ICT Express* 9 (3), 284–295. <http://dx.doi.org/10.1016/j.icte.2022.02.001>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S2405959522000200>.
- Van Loon, P., Deketele, L., Dewaele, J., McKinnon, A., Rutherford, C., 2015. A comparative analysis of carbon emissions from online retailing of fast moving consumer goods. *J. Clean. Prod.* 106, 478–486. <http://dx.doi.org/10.1016/j.jclepro.2014.06.060>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0959652614006489>.
- Wang, Z., He, S.Y., 2021. Impacts of food accessibility and built environment on on-demand food delivery usage. *Transp. Res. Part D: Transp. Environ.* 100, 103017. <http://dx.doi.org/10.1016/j.trd.2021.103017>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1361920921003151>.
- Xie, J., Xu, Y., Li, H., 2021. Environmental impact of express food delivery in China: the role of personal consumption choice. *Environ. Dev. Sustain.* 23 (6), 8234–8251. <http://dx.doi.org/10.1007/s10668-020-00961-1>, URL: <https://link.springer.com/10.1007/s10668-020-00961-1>.

- Yeo, V.C.S., Goh, S.-K., Rezaei, S., 2017. Consumer experiences, attitude and behavioral intention toward online food delivery (OFD) services. *J. Retail. Consum. Serv.* 35, 150–162. <http://dx.doi.org/10.1016/j.jretconser.2016.12.013>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0969698916302855>.
- Yuan, T., Da Rocha Neto, W., Rothenberg, C.E., Obraczka, K., Barakat, C., Turetti, T., 2022. Machine learning for next-generation intelligent transportation systems: A survey. *Trans. Emerg. Telecommun. Technol.* 33 (4), e4427. <http://dx.doi.org/10.1002/ett.4427>, URL: <https://onlinelibrary.wiley.com/doi/10.1002/ett.4427>.
- Zhang, S., Luan, H., Zhen, F., Kong, Y., Xi, G., 2023. Does online food delivery improve the equity of food accessibility? A case study of Nanjing, China. *J. Transp. Geogr.* 107, 103516. <http://dx.doi.org/10.1016/j.jtrangeo.2022.103516>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0966692322002393>.