Dynamic Parking Pricing Using Transaction Data in The City of Toronto

York University, Interactive-OR.net

Wenhan Luo, Hesam Rashidi, Fatemeh Sadeghi, Raneem Ayoub, Mehdi Nourinejad, Matthew J. Roorda

EXECUTIVE SUMMARY

Managing on-street parking in dense urban areas presents challenges due to high demand and limited supply. Dynamic parking pricing balances supply and demand by adjusting hourly rates based on observed occupancy levels of each parking zone. Parking rates are increased in high-occupancy zones and decreased in low-occupancy ones. Table 1 presents the summary of the findings and recommendations of this report. Successful pilots in cities like Washington DC, San Francisco, and Los Angeles have improved parking availability and lowered congestion through dynamic parking pricing. The parking rate adjustment cycle in such pilots has been any of one, three, six, or twelve months, thus deserving to term the technology “Adaptive Pricing” compared the “Dynamic Pricing,” as the latter can be misinterpreted for shorter adjust cycles such as a minute or hour.

Existing pilots have so far relied heavily on pavement-mounted sensors, magnetic pucks that detect the presence of a parked vehicle above them, and cameras to collect parking occupancy data for rate adjustment. These sensors, however, are infrastructure-invasive, costly to acquire and maintain, and hard to scale. In contrast, a highly neglected source of on-street parking occupancy data is payment transactions, which are digitally recorded and readily available.

This report has been commissioned by the Transportation Services Department of the City of Toronto to investigate the potential of parking transaction data in predicting occupancy levels for dynamic parking pricing. Processed transaction data from installed parking meters was used to predict occupancy levels based on historical data. Using a predictive algorithm removes the need to install occupancy detection infrastructure, reducing the overall costs of scaling dynamic parking pricing.

Analysis of parking transaction data revealed distinct demand patterns. On weekdays, there is a peak in the morning occupancy rates, likely due to commuters parking their vehicles in the parking facilities before heading to work. In contrast, the weekend parking patterns show that the afternoon occupancy rates surpass those of the morning, which could be attributed to leisure and social activities that are more prevalent during weekend afternoons. For nighttime occupancy, the rates present a consistent pattern across the entire week. The relatively uniform low occupancy during these hours may suggest that fewer individuals require on-street parking during the night, regardless of whether it is a weekday or weekend.

The proposed pricing algorithm uses a dynamic parking pricing system, optimizing parking price adjustments based on occupancy levels within a parking network. Opening

parking spaces for high-density parking areas and filling up parking spaces at low-density parking areas while also predicting the changes to nearby parking areas based on price adjustments. By analyzing transaction data, the algorithm detects parking zones where price adjustments are statistically likely to improve overall network occupancy levels. A Graph Neural Network (GNN) model predicts how the network would react to these price changes, predicting three variables: the average number of open parking spaces, average vehicle parking duration, and parking demand. The GNN model demonstrated high accuracy and generalization to new, previously unseen data.

The proposed pricing adjustments are expected to improve parking availability and balanced occupancy rates. For instance, in the morning period (7 AM – 12 PM), the price adjustments redistributed parking usage and made it easier to find spots in high-demand parking zones. The number of available spots increased by 8.5%, and the average occupancy rate decreased by 18.9% despite a 35.1% drop in parking demand. This occurs as the GNN predicted that vehicles will park for 5.6% longer durations. In the afternoon (12 PM – 5 PM), open spaces saw a modest increase of 1%, with occupancy rates rising by 4.6% due to a 25% increase in parking demand. During the evening period (5 PM – 9 PM), open spaces more than doubled, leading to a 3.6% reduction in occupancy rates, while parking demand surged by 57.1%. This seemingly large increase in demand is significant because the baseline demand in the evening is relatively low, so even a small absolute increase results in a substantial percentage change.

In addition to the pricing algorithm, an interactive web-based tool and ArcGIS dashboard were developed as part of this study to enhance parking management. The interactive web-based tool allows users to visualize various parking metrics on a map, such as occupancy, capacity, and pricing. This tool supports interactive manipulation, enabling users to adjust prices, capacities, and other parameters to simulate the impacts of these metrics on parking patterns. Furthermore, an auto-optimization feature calculates and implements optimal pricing adjustments across selected areas. Complementing this tool, an ArcGIS dashboard offers advanced data visualization and trend analysis, providing a comprehensive view of the parking network's status without optimization features. These tools facilitate effective decision-making and real-time management of urban parking resources.

Table 1 - Report summary, findings, and recommendations

Executive Summary	
Policy Problem	Managing on-street parking in dense urban areas due to high demand and limited supply.
Research Question	Can parking transaction data be used to predict occupancies for dynamic parking pricing effectively?
Key Findings	- Dynamic pricing has proven effective in pilot programs, balancing supply and demand. These pilots primarily used cameras and sensors which are infrastructure-intrusive, expensive to maintain, and hard to scale. - Payment transaction data is a viable alternative for determining occupancy rates as it is readily available in many cities and does not require additional investments. **Patterns observed in the data:** - Parking zones experience peak demand during weekday mornings likely due to commuters, where more than 50% of the zones have an occupancy rate of more than 20%. - During weekends, afternoon occupancy rates surpass those of the morning, which could be attributed to leisure and social activities. - Evening occupancy rates reflect that fewer individuals require parking during the night, regardless of whether it is a weekday or weekend. - Higher parking rates do not necessarily lead to lower occupancy; some high-priced zones maintain high occupancy due to the high demand and convenience of these locations.
Results	**Pricing adjustments using transaction data are expected to improve parking availability at congested zones and balance occupancy rates:** - Morning (7 AM – 12 PM): Reduced occupancy rates by 9.1%. - Afternoon (12 PM – 5 PM): 8.7% increase in the number of open spaces. - Evening (5 PM – 9 PM): Reduced occupancy rates by 6.7%.
Recommendations	- Implement a dynamic pricing algorithm based on transaction data. - Use predictive algorithms to reduce infrastructure costs. - Use web-based and ArcGIS tools for real-time management.
Conclusion	Predictive algorithms and transaction data can improve parking availability and reduce costs. Improved tools support effective decision-making and real-time management of urban parking resources.

TABLE OF CONTENTS

WHAT IS DYNAMIC PARKING PRICING?

Parking management is complex in dense urban areas due to the high demand for downtown cores and insufficient parking supply. Parking mismanagement can yield adverse consequences such as long parking search times and increased distances from destinations. Pricing is the conventional way of balancing supply and demand, and parking is no exception. Parking pricing strategies include fixed, hourly, progressive, time-of-day, and dynamic/adaptive. Fixed pricing charges drivers a fixed cost for a specified duration, often daily, making it suitable for commuting trips or long stays. Hourly pricing, widely applied, charges drivers proportional to their parking duration, making it suitable for short stays associated with shopping and recreational trips. Time-of-day pricing charges dissimilar rates during morning, afternoon, and evening hours. Progressive pricing charges a rate which increases if the parking duration surpasses a specified threshold (e.g., $3 per hour for the first hour and $5 per hour afterwards). Figure 1 demonstrates the total parking payment depending on the parking duration in each pricing strategy.

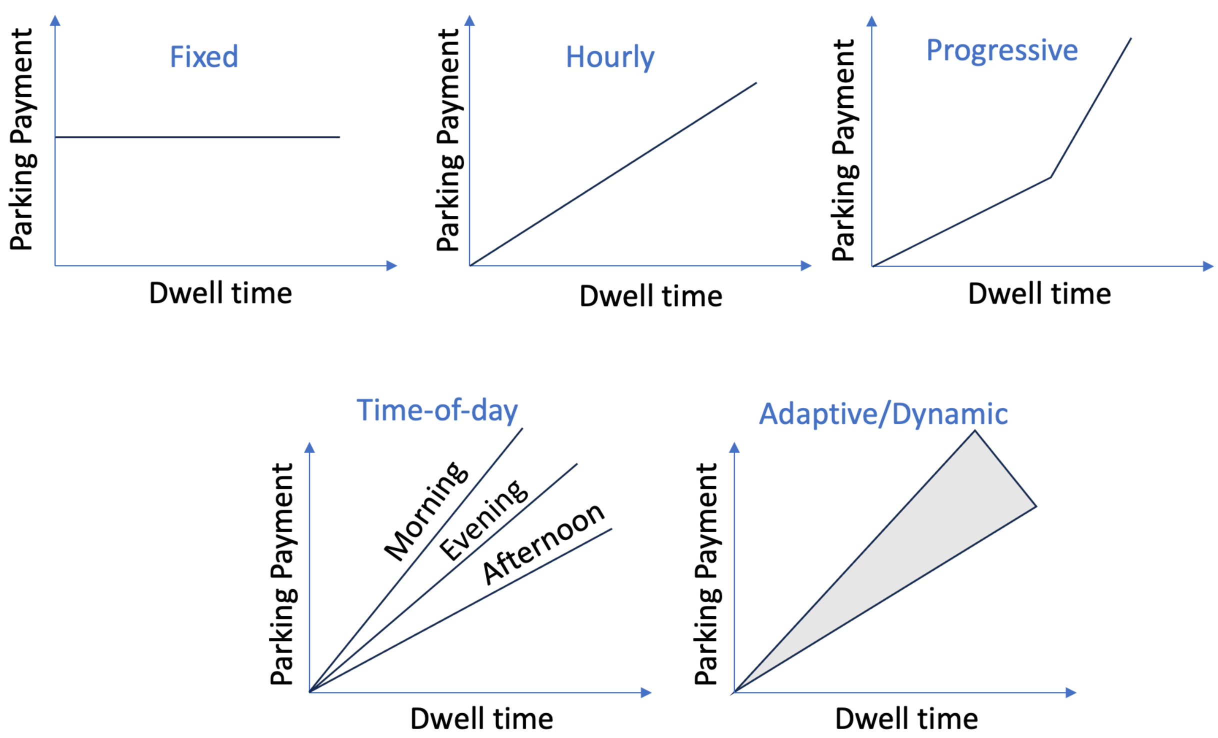

Figure 1 - Parking pricing strategies

Dynamic pricing is comparatively more advanced than other strategies as it allows for periodic adjustments (e.g., monthly, weekly, or even hourly) in the parking rates based on detected occupancy levels. Parking rates are adjusted to shift demand from high-occupancy to low-occupancy parking zones. Figure 2 (a) shows two nearby blocks priced equally, with the difference that the parking spaces in Street 1 are closer to attractive points of interest (e.g., coffee shops, restaurants, clothing stores), while those of Street 2 are not (e.g., commercial buildings). Therefore, Street 1 (with an occupancy of 100%) would likely have a higher occupancy than Street 2 (with an occupancy of 40%). Figure 2 (b) shows a conceptual implementation of dynamic pricing in the two streets, whereby the hourly rate of Street 1 increases from $3 to $4, and its occupancy ratio decreases as some drivers are deterred from parking or change their parking location. Likewise, as the parking price of street 2 drops from $3 to $2, street 2 attracts more drivers, increasing the occupancy ratio.

Figure 2- Dynamic parking pricing example.

PARKING OCCUPANCY DATA COLLECTION

Parking occupancy information is required for dynamic pricing because the price adjustments rely on the observed occupancies. There are several methods to acquire occupancy data, each with advantages and disadvantages. The main detection technologies are magnetic sensors, cameras, manual surveys, and transaction data.

Sensors

In-ground magnetic sensors are small-sized pucks that detect the presence of a single vehicle above them. Figure 3 (a) shows an example of a magnetic puck installed in the pavement using an adhesive. The pucks are self-powered wireless devices that send occupancy data to sensor management systems via a network of nearby gateways. SF-Park is a signature example of using magnetic sensors for a dynamic parking pricing pilot in San Francisco, with a total of 8,200 StreetSmart LLC in-ground sensors (SFPark, 2013).

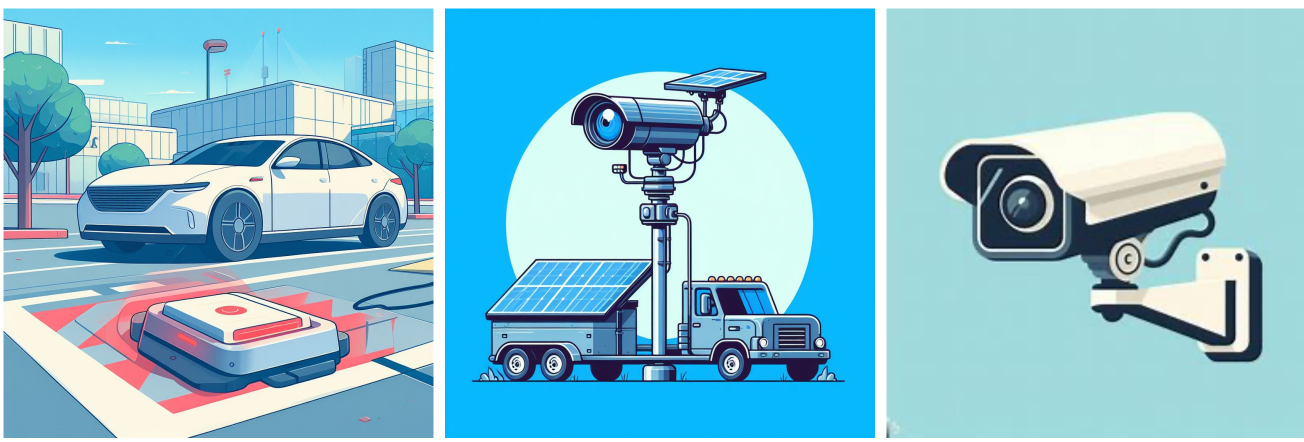

(a) Parking Sensor (b) Portable CCTV Camera (c) Fixed Camera

Figure 3 – Parking Occupancy Detection Methods

In the SF-Park pilot, the magnetic sensors encountered two main problems. First was electromagnetic interference caused by overhead lines for buses, light rail vehicles and alternating currents from junction boxes. Electromagnetic interference reduces the accuracy of the in-ground sensors, which will affect vehicle occupation detection. The second issue was early battery failures, as the sensors did not reach their 5-year battery life.

Since the deployment of magnetic sensors in SF-park, new sensor technologies are now available that use radar and optical readers and have an overall better battery life, averaging around 8-10 years before replacement. Newer sensor models offer adhesive options that are less intrusive to infrastructure. In a Washington DC pilot, Park-DC, multiple puck sensor vendors tested the effectiveness and limitations of their sensors. Due to magnetometer interference issues in past pilots, Park-DC used sensors with radar and optical functions to collect parking occupancy data. Park-DC also worked closely with the sensor vendors to identify and solve sensor communications issues (ParkDC, 2019) and adjusted the sensor ping frequency until it was at an acceptable rate to gather occupancy data.

Cameras

Cameras obtain parking occupancy information by detecting vehicles inside the camera's field of view. Three types of cameras can detect occupancy, each with advantages and disadvantages. The first type is portable CCTV Cameras, which capture the footage of a parking area, giving details on vehicle lengths, vehicle counts, vehicle speeds and vehicle classification. Portable CCTV Cameras are mounted on mobile trailers that can be relocated to record data at different locations, see Figure 3 (b). From portable CCTV cameras, the footage can be processed in two ways: either manually counting for small-scale data collection or using AI to determine the parking occupancy for large-scale data collection. While portable CCTVs offer certain advantages, their potential to block traffic and limited lens angles make them a less viable solution for measuring parking occupancy in urban areas compared to the alternative methods. In dense urban cities, the physical footprint of portable CCTVs can be a concern. These cameras can obstruct essential services or impede traffic when placed, despite their portable design. This issue is further compounded by their limited lens angles, which restrict the field of view and, thus, the effectiveness in monitoring a comprehensive area.

Time-lapse cameras offer an alternative method for occupancy detection (see Figure 3 (c)) by capturing images at specific intervals (e.g., five minutes). Park-DC (Washington DC Pilot) used time-lapse cameras at a frequency of 5 minutes. Time-lapse cameras are mounted on city assets, like streetlights or traffic signal poles and provide a wider lens angle, covering a larger area than a CCTV camera. Despite their potential for greater area coverage, time-lapse cameras are not the ideal solution for urban settings due to their high maintenance costs and the logistics and environmental challenges concerning their power source.

The third type is a fixed camera, which is mounted on streetlights or signal poles, similar to time-lapse cameras, and captures a narrow field of view. The main flaw of fixed cameras is the scarcity of feasible mounting locations that provide an acceptable aerial view, a communications line, and power. The installation requires a permit from the pole operators in charge.

Transaction Data

Access to occupancy information requires a cost-effective, practical, and scalable approach. An often disregarded but widely available data source that can be used for parking occupancy prediction is the parking transaction data, which consists of information about each parking location's date and time of transaction and the price paid. Parking transaction data is readily available in many cities and does not require investment in city infrastructure which can be costly and timely.

The main challenge with transaction data is that it does not capture permit-holders, illegal parkers (i.e., those who do not pay at all and those who overstay their payment), and early departers (i.e., those who leave earlier than granted). Parking occupancy inference from transaction data decreases in accuracy with more illegal parking and permit holders. Nevertheless, some pilots show reductions in illegal parking when dynamic pricing is implemented since the drivers would have better chances of finding a legal parking space. The Boston Sea Port pilot (City of Boston, 2016) reported noticeable reductions in double parking after implementing dynamic parking pricing. Due to higher parking availability, drivers are more willing to park legally as vacant parking spots are now open (City of Boston, 2016).

Hybrid Sources and Data Fusion

Data fusion combines multiple data sources, including magnetic pucks, cameras, and transaction data, to predict parking occupancy. Park-DC implemented data fusion by merging transaction data, pay-by-phone transactions, spatial occupancy sampling and parking citations (ParkDC, 2019). Figure 4 shows a combination of real-time and historical data sources that detect real-time occupancy and availability.

The Park-DC pilot comprised three phases to predict occupancy using data fusion. The first was using temporal data collected from portable CCTV to determine where to place in-ground sensors. The portable nature of the cameras allows them to be relocated to collect data from multiple sites and identify parking locations where demand patterns are unusual, potentially requiring an in-ground sensor to gather occupancy data. The second phase was to refine the sensor occupancy information using data fusion. The combination of real-time and historical occupancy data improved occupancy predictions for dynamic parking pricing. The third phase involved finding the minimal viable sensor coverage. In Park-DC's case, 50% sensor coverage was initially assumed. The city then randomly removed specific data sources to determine if occupancy could still be accurately predicted. The pilot aimed to gauge the importance of each data source for occupancy prediction with this approach. After the testing, Park-DC found that 50% sensor coverage was viable enough for the pilot area.

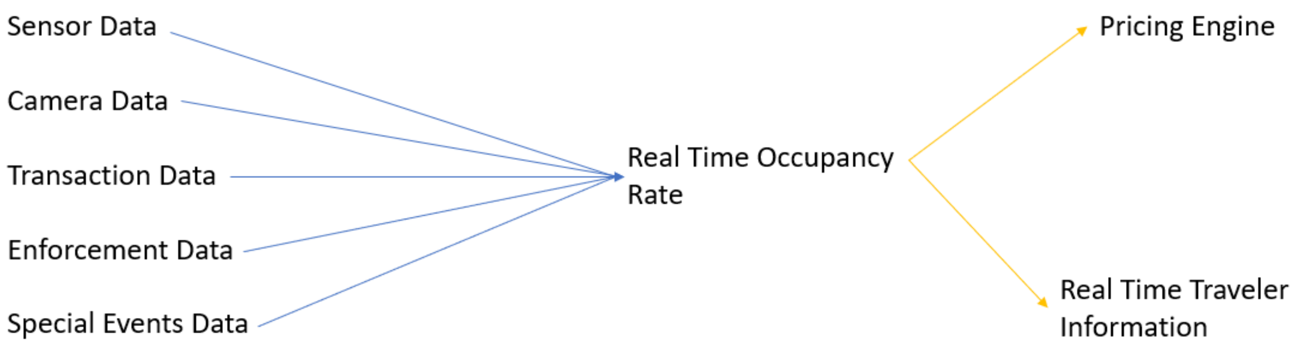

Figure 4 - Data Fusion Diagram

DYNAMIC PARKING PRICING PILOTS

City Pilots Overview

Table 2 presents the dynamic parking pricing strategies of select cities, highlighting how these cities have customized their strategies to fit local conditions. Despite shared goals in pricing and occupancy rates, the variations in the strategies highlight the lack of a universal solution. The following section reviews the most relevant pilots of dynamic parking pricing. The reviewed pilots are existing cities that have implemented a variation of dynamic parking pricing. San Francisco, City of Boston, and Washington DC are the investigated pilots as they are the only cities that generated comprehensive reports on their dynamic parking pricing pilots.

San Francisco Pilot (SF-Park)

San Francisco launched SF-Park, a pilot that implemented dynamic parking pricing to better balance the parking supply and demand. About one-quarter of San Francisco's metered parking spaces were monitored with magnetic detection pucks. The parking rates were increased if the occupancy was too high (>85%) and decreased if the occupancy was too low (<60%), intending to keep at least one parking spot open on every block.

As an outcome of the pilot, the chances of experiencing occupancies between 60% to 80% increased by 31%. Moreover, the average park search time decreased by 43%, and vehicle kilometers travelled reduced by 30% (SFPark, 2013). Some parking meters were upgraded to allow for dynamic parking price changes, with the technology to implement frequent changes to rates remotely, needless of a technician to change them manually.

Upgraded parking meters added more functionality to the existing parking system. Real-time occupancy data was disseminated through web and mobile apps, using the aggregated information to help drivers make informed decisions about their travel plans. Upon implementing dynamic pricing, SF-Park installed roadway sensors to evaluate the traffic impacts of the initiative. The sensors gauged traffic volumes and speeds, allowing the city to appraise the impact of dynamic parking pricing on vehicular flow. The assessment helped determine whether the pilot program was beneficial in reducing traffic congestion.

Table 2- Dynamic Parking Pricing City Overview

City	Occupancy Detection Method	Time-of-day price variations	Average Period Taken for Price Change	Year Conducted	Target Occupancy
Seattle (City of Seattle, 2017)	Mix of Surveying, Transaction Data, and Historical Occupancy Data	Three Weekday Periods	Annually	2010	70-85%
San Francisco (SF-Park)	In-Ground Sensors	Three Weekday Periods	3 Months	2011	60-80%
Los Angeles (LAExpressPark)	In-Ground Sensors	Three Weekday Periods	3-6 Months	2012	70-90%
Calgary (City of Calgary, 2022)	Unknown	Four Weekday Periods	Annually	2013	50-80%
Washington DC (Park-DC)	Variety of methods, mostly sensors	Three Weekday Periods	3 Months	2016	None Stated
Boston (Performance Parking Pilot)	In-Ground Sensors	Three Weekday Periods	2 Months	2017	60-80%
Baltimore (Demand-Based Parking Meter Rate Setting) (City of Baltimore, 2017)	Surveys	None	6 Months	2017	75-85%

Boston Pilot

In 2017, the City of Boston initiated a dynamic parking pricing pilot in the Seaport area, mirroring the objectives of the SF-Park pilot in San Francisco. The primary goal was to maintain a target occupancy rate between 60 – 80%, thereby encouraging drivers to either park in less congested areas or opt for alternative transportation methods. The pilot used sensors mounted on parking meters to monitor parking occupancy. These sensors informed pricing adjustments: rates would increase by 50 cents if occupancy surpassed 80%, with a capped hourly rate ranging between $1 and $4.

Concurrently, Boston explored new mobility strategies to reduce reliance on car parking. The initiatives included allocating more spaces for car-share operators and expanding the bike-share network. These measures were aimed at promoting alternative transportation modes. The pilot’s impact assessment revealed notable reductions in illegal parking—12% in Back Bay and 35% in Seaport. Notably, the parking facilities rarely reached full capacity, which implies a higher availability of spaces. However, the pilot's effectiveness in Seaport was less apparent due to ongoing construction, new developments, and fluctuating seasonal demands, leading to minimal congestion reduction relative to operational costs.

Boston's pilot underscored several insights. Foremost was the importance of effectively communicating parking price changes to drivers, enabling better trip planning, and reducing the likelihood of cruising for parking. Additionally, the city identified limitations in the current sensor technology. The sensors provided unreliable occupancy data, thus challenging the scalability of this approach and complicating the assessment of demand levels. Consequently, Boston's experience suggested that alternative methods for measuring parking occupancy should be considered.

Washington DC Pilot (Park-DC)

In response to increasing traffic and limited parking, Washington D.C. launched Park-DC. Park-DC had three primary objectives: reducing parking search times, alleviating congestion, and developing effective parking management solutions with minimal investment in new assets. Park-DC's core strategy was to implement dynamic pricing, adjusted by demand and time of day. For occupancy detection, the city used a combination of parking sensors and surveillance cameras, supplemented by data from parking transactions. In addition to technological solutions, Park-DC focused on effective communication with drivers by installing regulatory signs to inform drivers about price changes and parking regulations. The effectiveness of the signs was assessed through surveys, providing insights into whether further improvements were necessary.

Park-DC also developed a mobile application to assist drivers with parking. The app facilitated the parking process and enabled the collection of valuable data on parking sessions, driver behaviour, and system performance. To evaluate the pilot's impact on congestion reduction and safety, Park-DC used metrics from manual surveys, parking citations, Bluetooth sensor data, census information, and camera footage. Finally, Park-DC explored the scalability of different occupancy detection methods for city-wide implementation, focusing on cost-effective solutions.

PRICE ADJUSTMENT

Dynamic pricing strategies typically divide the day into distinct periods that correspond with changes in parking demand (e.g., morning, afternoon, and evening). The adjustment cycle – the timing of these changes – is set to be frequent enough to respond to demand without causing confusion for drivers. The amount of rate change influences parking choices: too small a change might not influence the driver's behaviour, while a substantial change could discourage their decision to park in a parking zone. Data-driven insights inform these adjustments, ensuring that decisions are grounded in actual occupancy patterns.

Time of Day Pricing

Dynamic pricing for parking considers the variability of demand throughout the day, typically segmenting it into three periods. Mornings (e.g., 7/8 am to 11 am) are priced to accommodate the work commute rush, afternoons (e.g., 11 am to 4/5 pm) offer lower rates due to reduced parking needs, and evenings (e.g., 4/5 pm to 8/9 pm) adjust rates for those returning home. The parking demand is considerably lower and less flocculating between 8/9 pm to 7/8 am, hence, dynamic pricing is not advised for the said period. The mentioned tripartite structure, as shown in Table 2, is a model adopted by pilots such as SF-Park, LAExpressPark, Park-DC, and Boston's Performance Parking Pilot, aiming to match pricing more closely with the daily parking demand.

Adjustment Cycle

The frequency for adjusting parking rates must be carefully calibrated to avoid overwhelming drivers with excessive real-time changes. Park-DC stated that well-communicated pricing changes are more influential than frequent price changes (ParkDC, 2019). This approach facilitates better preparation for consumers in anticipation of price adjustments while allowing municipalities a broader window to ascertain more practical pricing schemes. As shown in Table 3, the median interval for price adjustments is approximately four months.

Rate Changes

The impact of parking price adjustments on user behaviour is complex and varies by location. Park-DC and SF-Park pilot programs used data-driven adjustments rather than solely relying on theoretical price change predictions. Park-DC took a rate price band approach to determine the prices of parking areas. It started using three price bands of $2.00, $2.30, and $2.75, but in the last phase, the price band increased to $10. The program expanded its pricing strategy to ten bands in the final phase, as shown in Table 3. The rationale for this increase was to leverage a more assertive pricing scheme significant enough to influence users' trip-planning behaviour. In contrast, SF-Park used a rule-based approach to pricing adjustments. This method entailed a $0.25 increase per hour for occupancy levels between 80-100%, no change for 60-80% occupancy, a $0.25 decrease for 30-60% occupancy, and a $0.50 reduction when occupancy fell below 30%. While simpler, this model was less aggressive in altering user behaviour. The primary objective of SF-Park was to alleviate congestion and optimize the use of parking spaces, even if it meant forgoing some revenue. As an initial pilot, SF-Park's conservative approach aimed to improve social welfare. Both pilots took a different approach as to how price adjustments are made. Still, both cities had a common strategy of testing pricing changes during the pilot to ensure parking price changes were data-driven and not through prediction of occupancy rate from a price change.

Table 3 - ParkDC Pilot Rate Structure (ParkDC, 2019)

Price Change	Rate Structure (Hourly Rates)
Baseline	$2.30
Round 1 October 2016	$2.00 $2.30 $2.75
Round 2 February 2017	$1.50 $2.00 $2.30 $2.75 $3.25
Round 3 May 2017	$1.00 $1.50 $2.00 $2.30 $2.75 $3.25 $4.00
Round 4 August 2017	$1.00 $1.50 $2.00 $2.30 $2.75 $3.25 $4.00 $4.75
Round 5 November 2017	$1.00 $1.50 $2.00 $2.30 $2.75 $3.25 $4.00 $4.75 $5.50

USER COMMUNICATION METHODS

Two important challenges are to be addressed in any application of dynamic parking pricing: first is the method of parking occupancy detection and second is the information dissemination method to inform drivers of the changes in the price rates. Drivers are to receive rate adjustments in a timely manner to plan their trips accordingly and make parking arrangements. This section discusses the prominent strategies to foster such effective communication.

Web/Phone Applications

Cellphones and web platforms are reliable and accessible methods for communicating rate adjustments. Figure 5 shows LAExpressPark's method of displaying parking information on their website. The interface allows users to select a specific area and peruse the available pricing options within that chosen region. Additionally, the platform offers the flexibility of adjusting prices and parking durations to accommodate the specific requirements of each driver. In contrast, based on Figure 6, the phone application created by Park-DC shares essential information on the parking zones, providing parking rates at each time-of-day and other important parking spot information. The app provides all the necessary information for users to plan their trips.

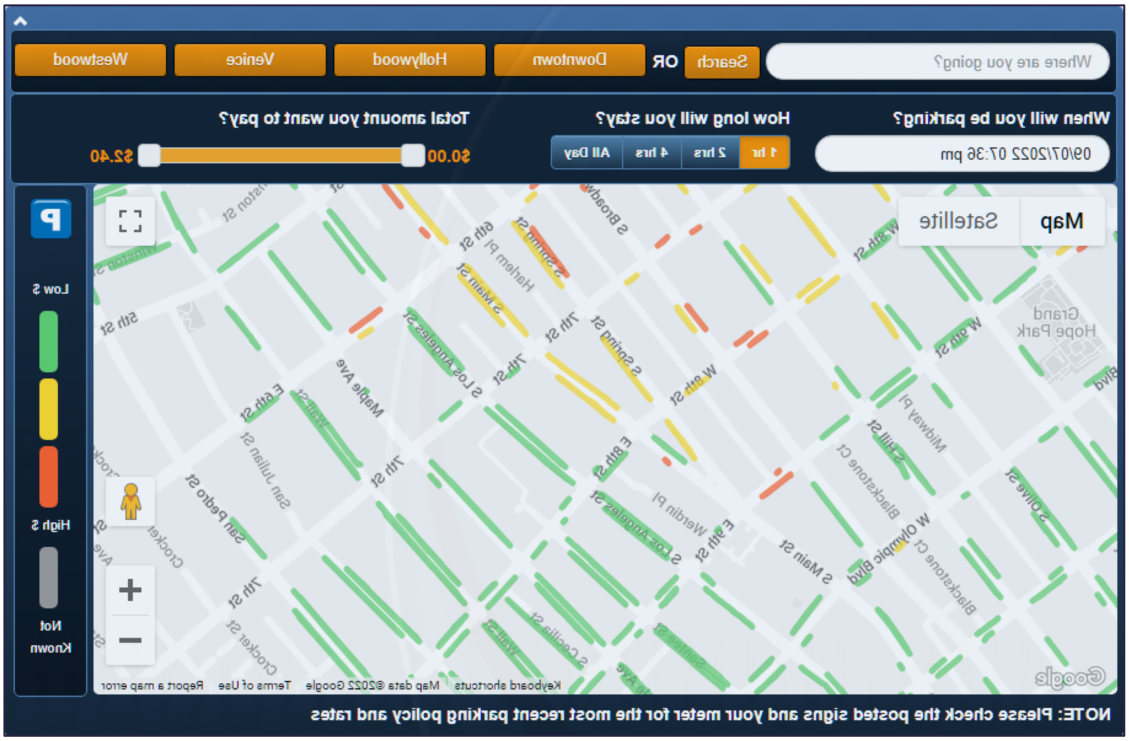

Figure 5 - LAExpressPark Website (LA Express Park, 2012)

Web platforms offer the advantage of rapid updates and widespread accessibility. However, they are not without their limitations, particularly concerning drivers who either lack access to a computer or smartphone or face challenges in navigating these digital tools.

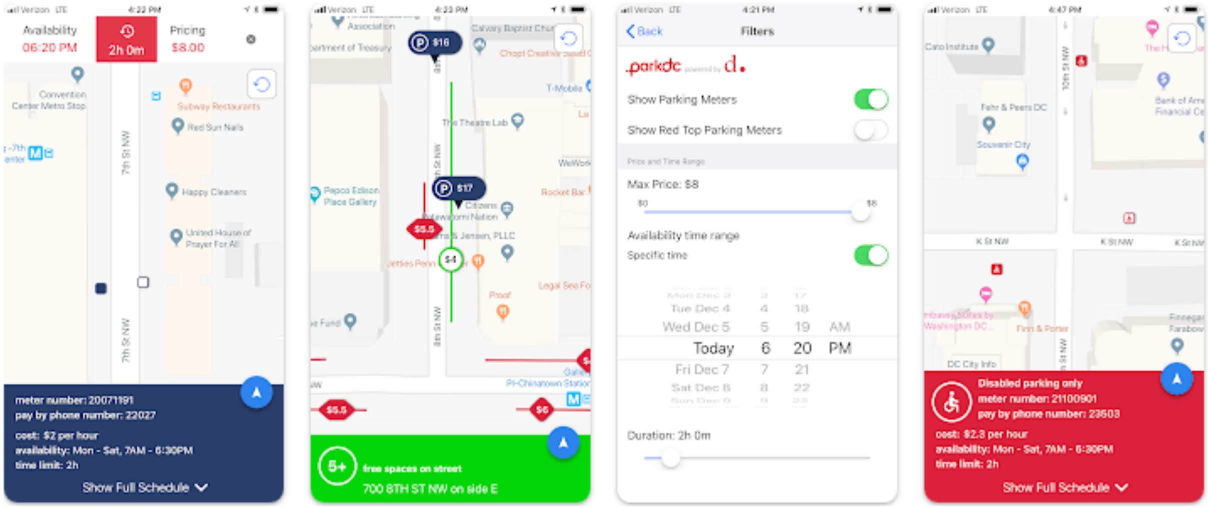

Figure 6 - Park-DC Phone App (ParkDC, 2019)

Parking Guidance Signs

Effective communication of real-time parking space availability is essential for guiding drivers toward these less congested areas. While websites and mobile applications are useful tools for pre-trip planning, they may not be universally accessible to all drivers. Physical, on-site parking guidance signs could bridge this gap and may be particularly beneficial for those already en route and seeking parking.

There are two parking guidance signs: variable message signs and wayfinding signs. Variable message signs are more costly but also efficient in directing traffic to available parking spaces. Digital signs can project parking locations live, as they receive information from a network with live occupancy rates. As traffic changes throughout the day, the sign can change with it to direct traffic effectively. Figure 7 shows an example of a variable message sign used in San Francisco that showed the number of parking spaces in garages in the vicinity.

The other type of physical, on-site sign, wayfinding signs, is meant to guide drivers to specific parking locations, so they can easily reach their destination. These fixed signs are mounted on poles near all parking areas. Fixed signs are less costly and effective in guiding drivers to find parking in less congested spots. Figure 8 shows a San Francisco wayfinding sign used during the pilot project. A simple direction and name of the location are on the sign to aid drivers in navigation.

Figure 7- SF-Park Variable Message Signs (SFPark, 2013)

Figure 8 – SF-Park wayfinding sign (SFPark, 2013)

PARKING TRANSACTION DATA FROM TORONTO

This section presents a detailed analysis of the transaction data from the Toronto Parking Authority (TPA) to identify trends and patterns in parking behaviour and adjust parking rates using dynamic pricing. The TPA, which oversees most parking facilities in the Greater Toronto Area (GTA), provides a comprehensive data set of parking transactions. However, parking transaction data only includes paid parking and does not capture instances of illegal parking or usage by permit holders.

Dataset Description

The dataset contains raw transaction data for approximately 1000 on-street parking areas throughout the GTA, as shown in Figure 9. The data dates range from May 2019 to July 2019 and May 2022 to July 2022. Table 4 describes each column in the data set.

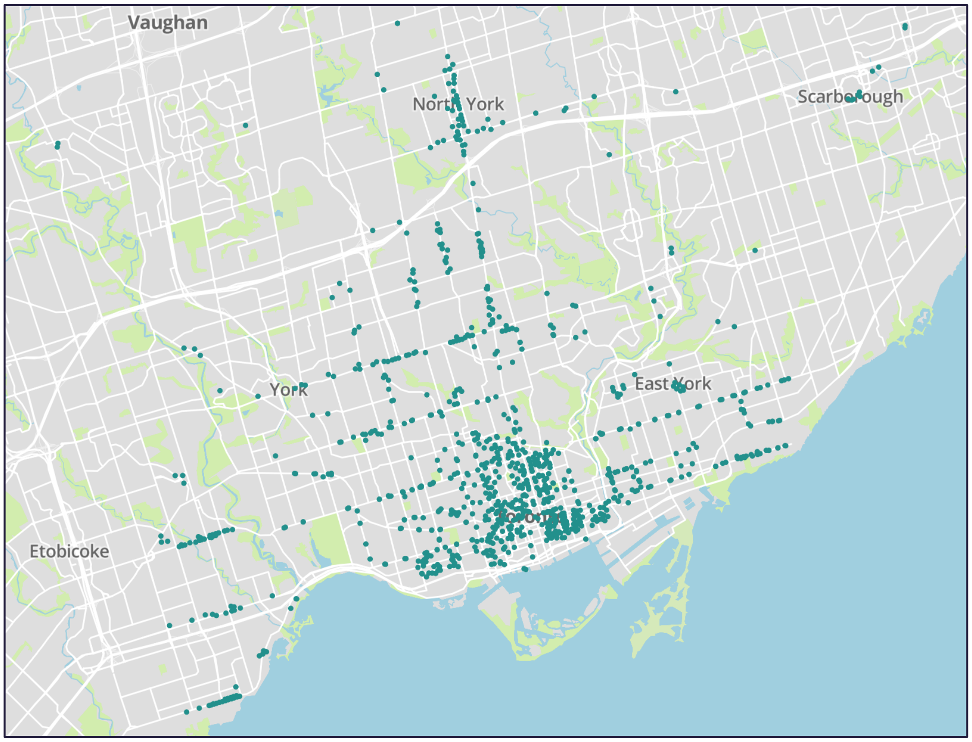

Figure 9 – Parking locations in the Toronto Parking Authority transaction dataset.

Occupancy Graph Format

Figure 10 - Figure 12 exhibit occupancy comparisons as of May 1st, 2019, for a sample parking zone (i.e., location ID 3001, situated on Elm Street between Yonge St. and Bay St). On each graph, the horizontal axis indicates time by the hour, while the vertical represents the occupancy count, reflecting the number of parked vehicles at given times. The red line across the graphs marks the total available parking capacity. Each of the three graphs differs by the time interval of occupancy counting: 1 minute, 15 minutes, and 30 minutes. These intervals affect the frequency of data capture for vehicles entering or leaving the parking area. Shorter intervals yield high-resolution data but can clutter the graph, obscuring broader trends. Conversely, longer intervals smoothen the graph but may omit finer details. A 15-minute counting interval is selected for subsequent data analysis as it provides a balance, offering sufficient detail for trend interpretation without overcomplicating the visual presentation.

Table 4 - Dataset Descriptions

Variable Name	Column Description
Location ID	Distinct location identifier for each parking zone.
Operation Type	Determines the operation type of the parking area (i.e., on- or off-street).
Operation Owner	The ownership entity of the parking area.
Parking Meter Code	Identifies the specific meter associated with each transaction. Transactions conducted via mobile payment are assigned a meter code that matches the location ID of the parking area.
Transaction Start Time	Transaction start time The outline for a transaction time is in the format of YYYY-MM-DD hh:mm:ss.
Transaction End Time	Transaction end time of the parking, formatted as YYYY-MM-DD hh:mm:ss. Note that this indicates the time paid for, not necessarily the time when the user departs from the parking area.
Amount Paid	Indicates the payment amount for each transaction, reflecting the sum paid by the user either at the parking meter or via the mobile app.
Transaction Type	There are two types of transactions available: payments made at the parking meter in the parking area, and payments processed through the TPA mobile app for a parking spot.
Parking Duration	The parking duration paid for by the user, measured in seconds.

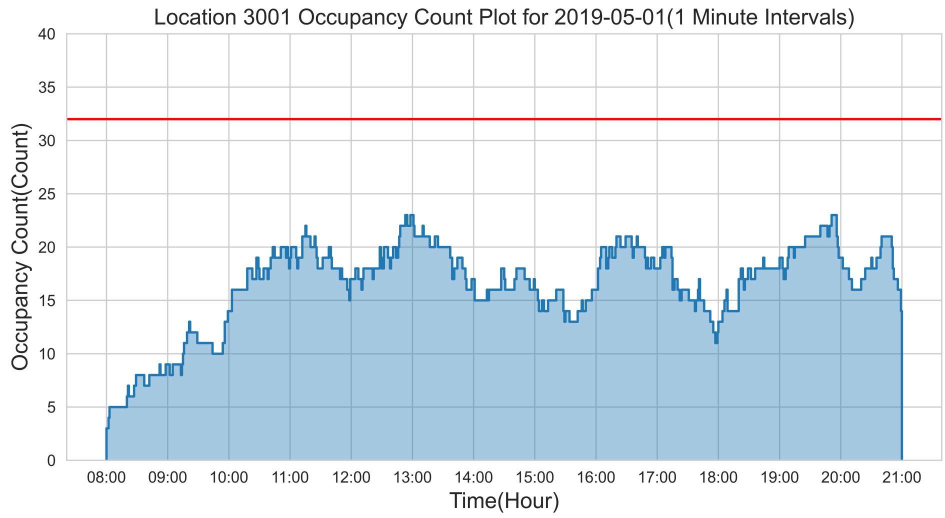

Figure 10 - Occupancy Count Plot (1-minute Intervals). The red line demonstrates the parking capacity.

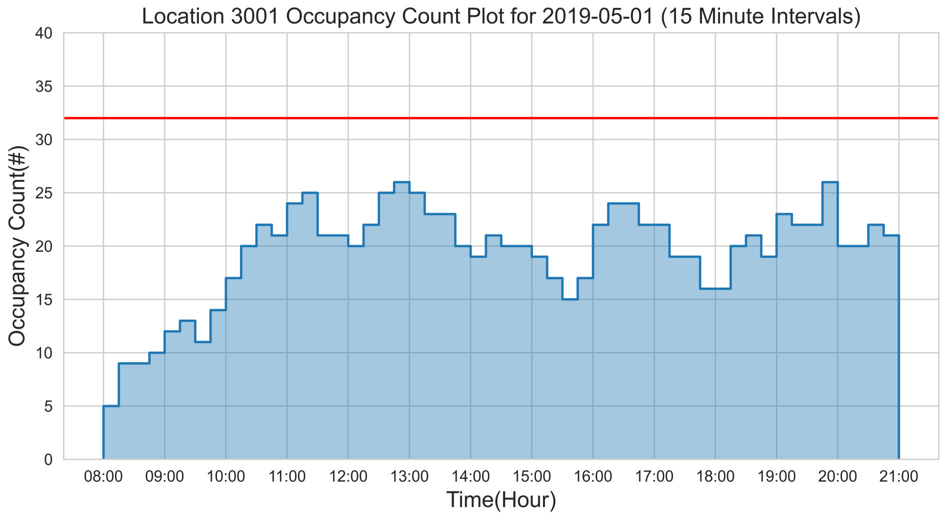

Figure 11 - Occupancy Count Plot (15-minute Intervals). The red line demonstrates the parking capacity.

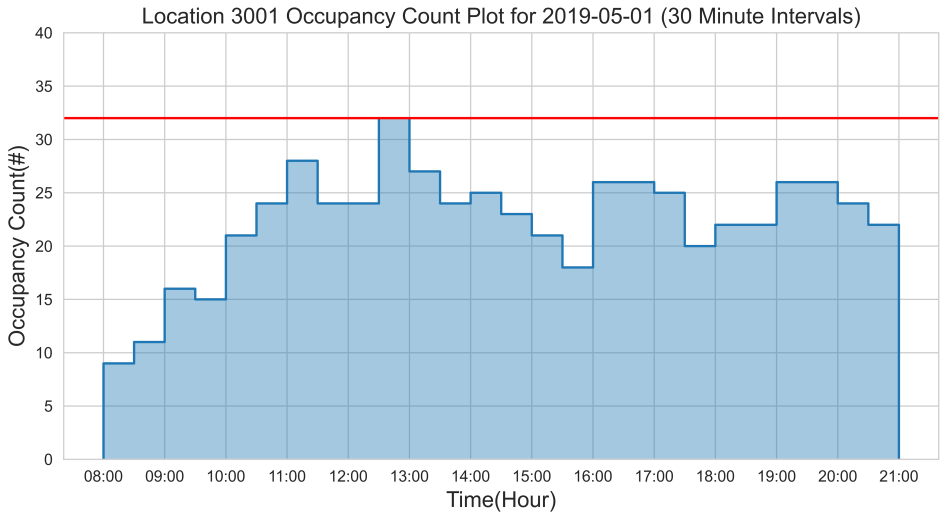

Figure 12 - Occupancy Count Plot (30-Minute Intervals). The red line demonstrates the parking capacity.

Day-of-Week Analysis

This section of this report analyzes the average occupancy rate trends for all parking locations, comparing weekdays to weekends. As depicted in Figure 13 - Figure 17, on weekends, occupancy has the higher average occupancy rates, with notable influxes in the afternoon and evening. Saturday's experience peaks in the afternoon and evening, reflecting less demand in the morning. Sundays show sustained occupancy into the afternoon and early evening, but Sunday hours of allowed parking range from 11:00 am to 9:00 pm for most parking locations. Conversely, on weekdays, there are clear occupancy peaks in the morning and a consistent rise in the afternoon and evening, indicating regular entry of drivers into the parking area. These observations suggest the importance of implementing time-of-day pricing variations to manage parking demand effectively.

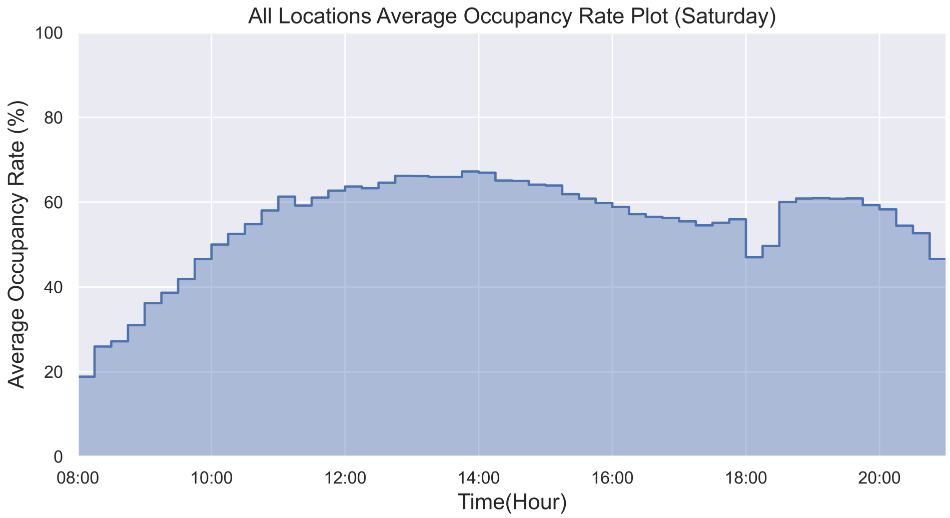

Figure 13 - Average Occupancy Rate Plot for All Locations (Saturday)

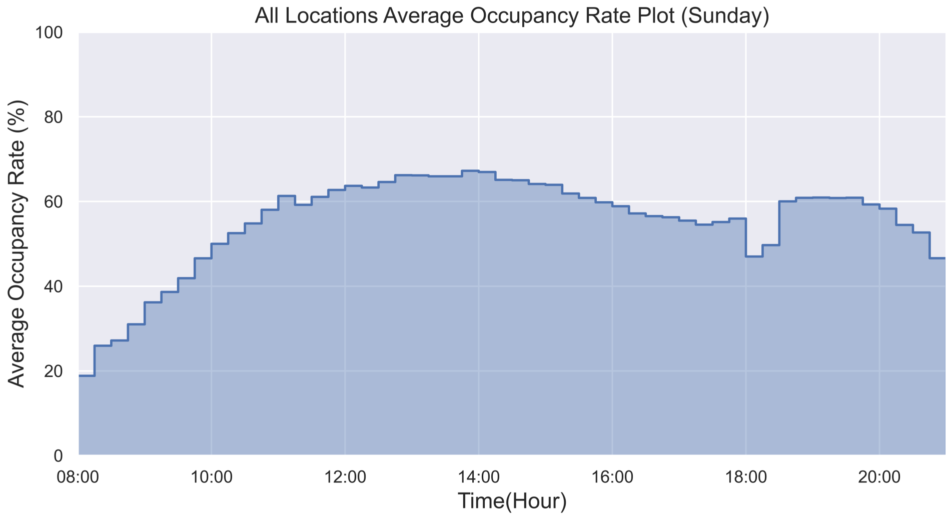

Figure 14 - Average Occupancy Rate Plot for All Locations (Sunday)

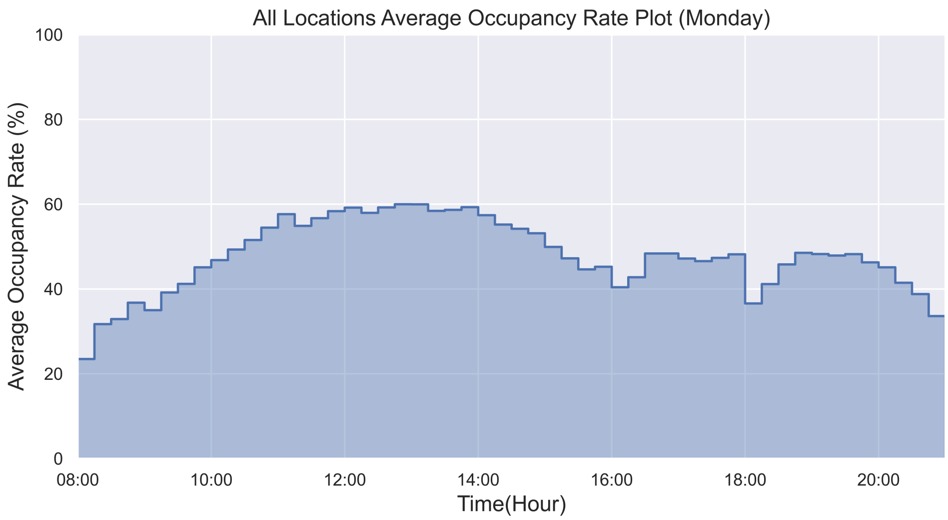

Figure 15 - Average Occupancy Rate Plot for All Locations (Monday)

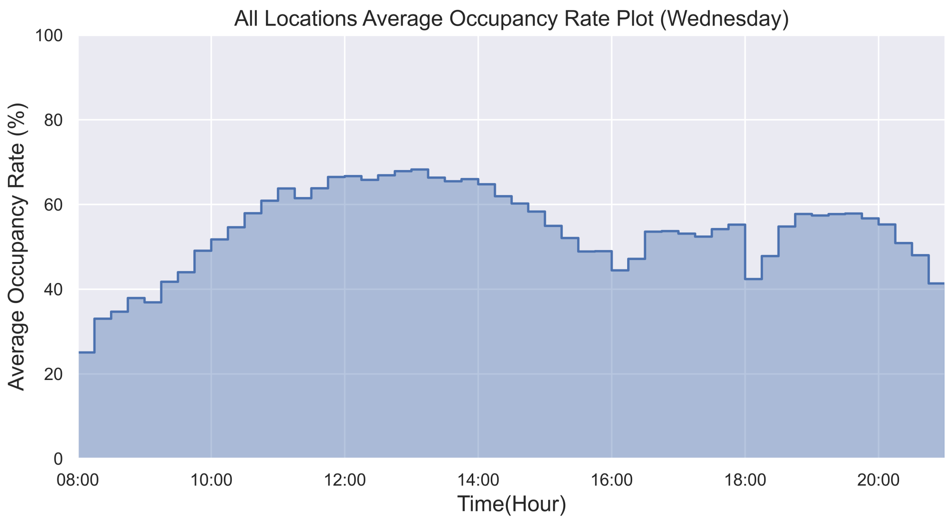

Figure 16 - Average Occupancy Rate Plot for All Locations (Wednesday)

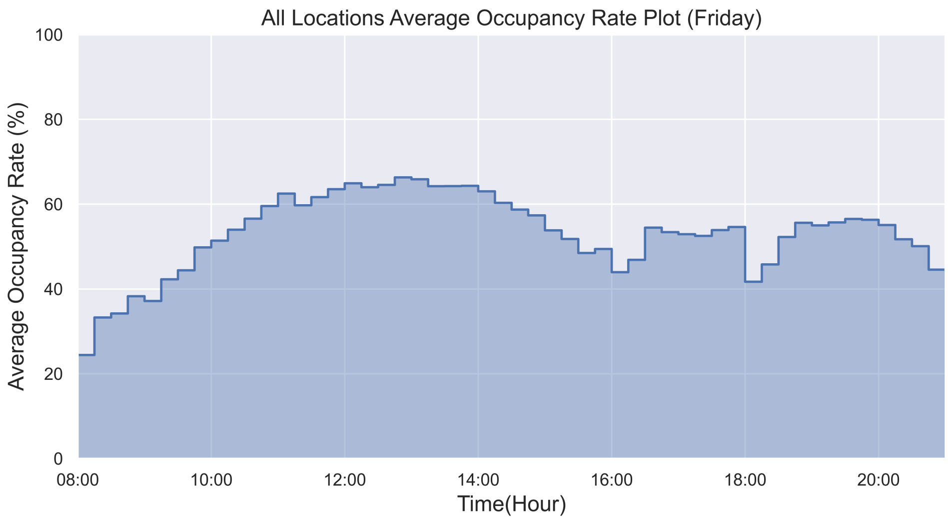

Figure 17 – Average Occupancy Rate Plot for All Locations (Friday)

Daily and Average Monthly Comparisons

Figure 18 illustrates the average occupancy rates for parking locations, segmented into three intervals (referred to as the time of days - TODs): morning (i.e., 7 am -12 pm), afternoon (i.e., 12 pm - 5 pm), and night (i.e., 5 pm – 9 pm). These rates are calculated by dividing the average number of occupied spots during a given interval by the parking facility's total capacity. As shown in Figure 18, there is a peak in the morning occupancy rates, which aligns with conventional workday routines, where individuals tend to occupy parking spots during morning hours for employment purposes. The data suggests that typically more than 50% of the parking facilities have an occupancy rate of more than 20% during weekday mornings, likely due to commuters parking their vehicles before heading to work.

The weekends show a notably different pattern. The afternoon occupancy rates surpass those of the morning, which could be attributed to leisure and social activities that are more prevalent during weekend afternoons. For nighttime occupancy, the rates present a consistent pattern across the entire week. The uniformly low occupancy during these hours may reflect a common behaviour pattern, where fewer individuals require parking during the night, regardless of whether it is a weekday or weekend.

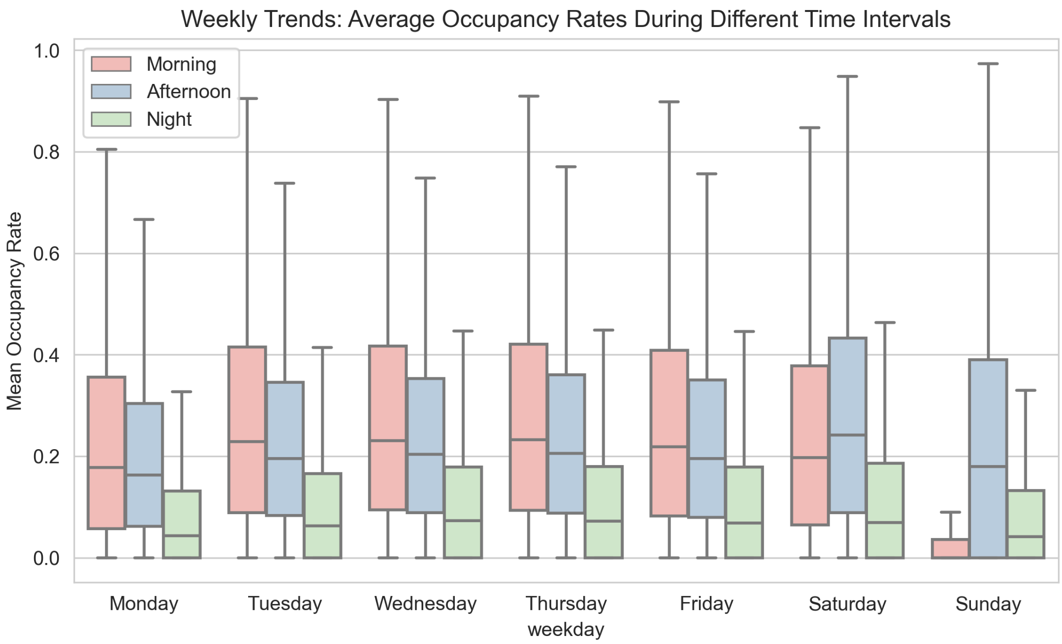

Figure 18 - Parking patterns between weekdays and weekends, with work-related activities driving higher morning occupancy during the week and social events influencing weekend afternoons.

Spatial Pricing Patterns

The transaction data suggests that higher rates do not always lead to lower occupancy. In fact, as Figure 19 - Figure 21 show parking zones with higher hourly prices typically have a larger occupancy rate. This observation specifies a spatial dimension in parking behaviour where areas with higher demand can command and sustain higher prices. Figure 22 displays the distribution of hourly parking rates by zone, showing that the highest rates are persistently found in the downtown core.

Figure 19 - Morning occupancy rates across hourly rates

Figure 20 - Afternoon occupancy rates across hourly rates

Figure 21 - Evening occupancy rates across hourly rates

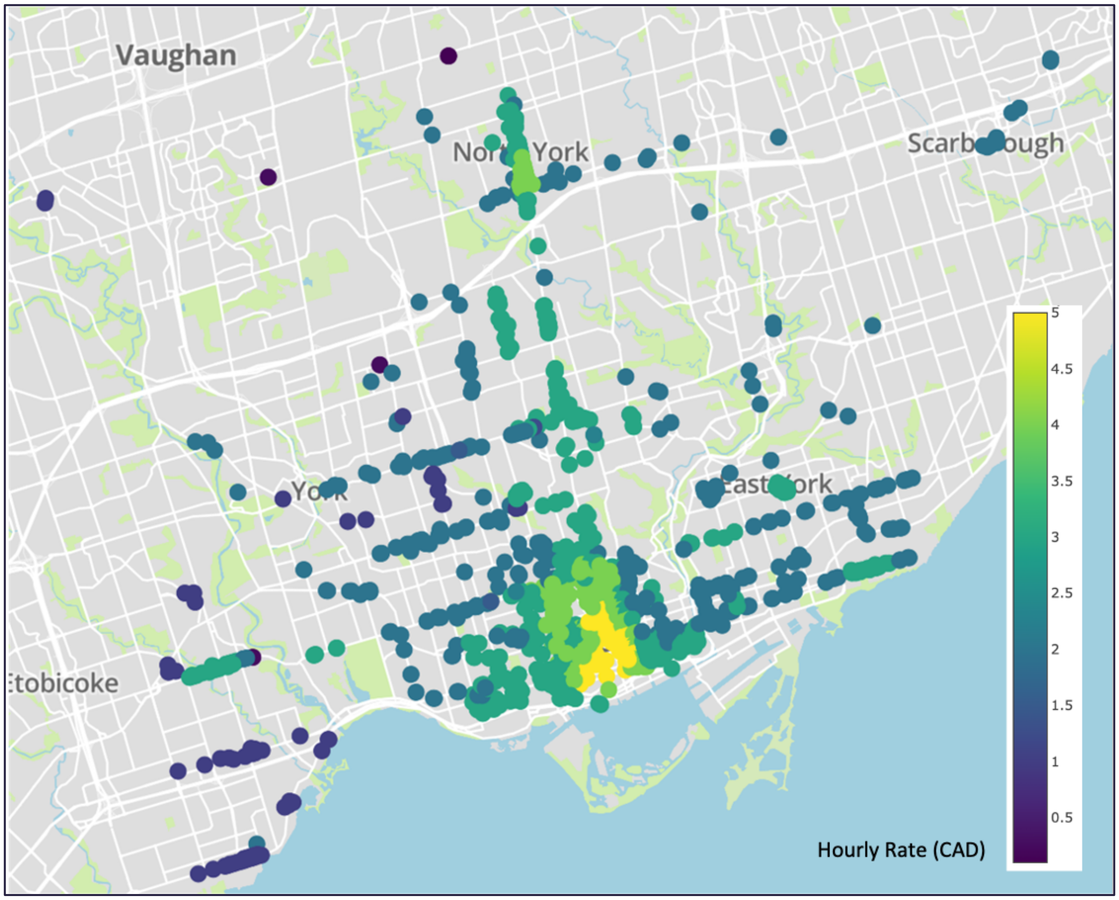

Figure 22 - Varying hourly parking rates across different locations. There is a trend for locations situated closer to the downtown core to have higher prices.

Proximity Analysis

The transaction data shows the efficacy of dynamic pricing in parking management. For instance, Figure 23 and Figure 24 highlight two contrasting occupancy scenarios in closely situated parking zones. In Figure 23, the parking zones on Toronto St (Location ID 4111) and The Esplanade (Location ID 4131) are seven minutes apart by walking, and their occupancy rates are similar, while a $1 discrepancy in parking rates exists. Both locations maintain occupancy within the ideal range of 60-80%, suggesting that current pricing is effective and does not necessitate adjustments. Conversely, Figure 24 shows two adjacent parking areas on Wellington St, locations ID 4120 and 4121. Despite their close vicinity, there is a marked disparity in their occupancy rates. Location ID 4120, with a higher rate of $5/hr, seldom achieves 20% occupancy, whereas Location ID 4121, at $4/hr, frequently experiences high occupancy. This disparity indicates a clear opportunity for dynamic pricing: reducing rates at location ID 4120 or increasing rates at location ID 4121 could balance occupancy levels more effectively.

<table>
<tr><th>Street Comparison</th><th>Map Image</th></tr>
<tr><td>Toronto St. (4111):

• 25 Spaces
• $5.00/hr
• Average Occupancy = 46.85%

The Esplanade (4131):

• 29 Spaces
• $4.00/hr
• Average Occupancy = 49.38%</td><td>
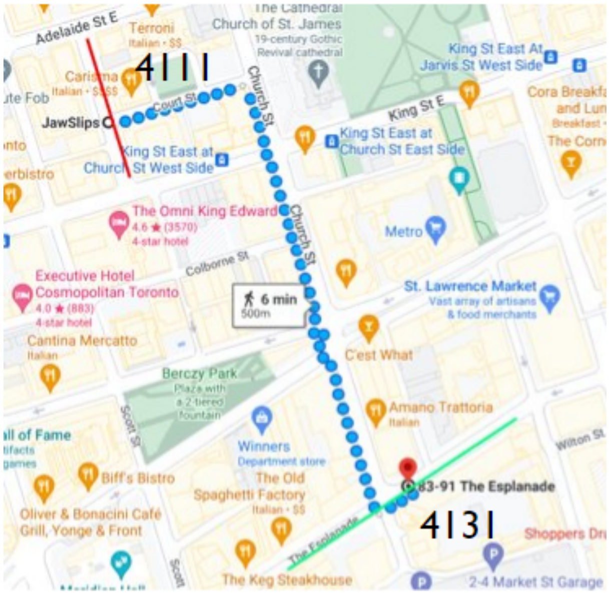

</td></tr>
</table>

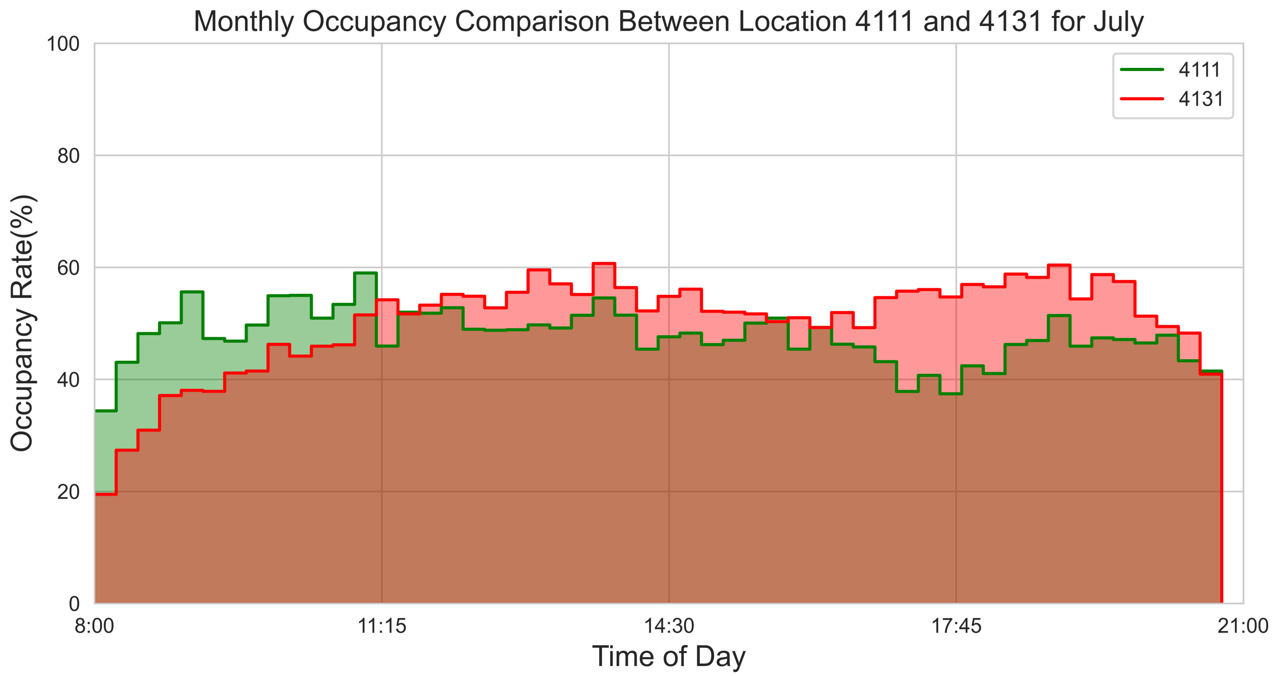

Figure 23 - Comparison between Toronto St and The Esplanade

Street Comparison	Map Image
WELLINGTON ST E (4121): • 14 Spaces • $4.00/hr • Average Occupancy = 66.18% WELLINGTON ST E (4120): • 27 Spaces • $5.00/hr • Average Occupancy = 11.51%	

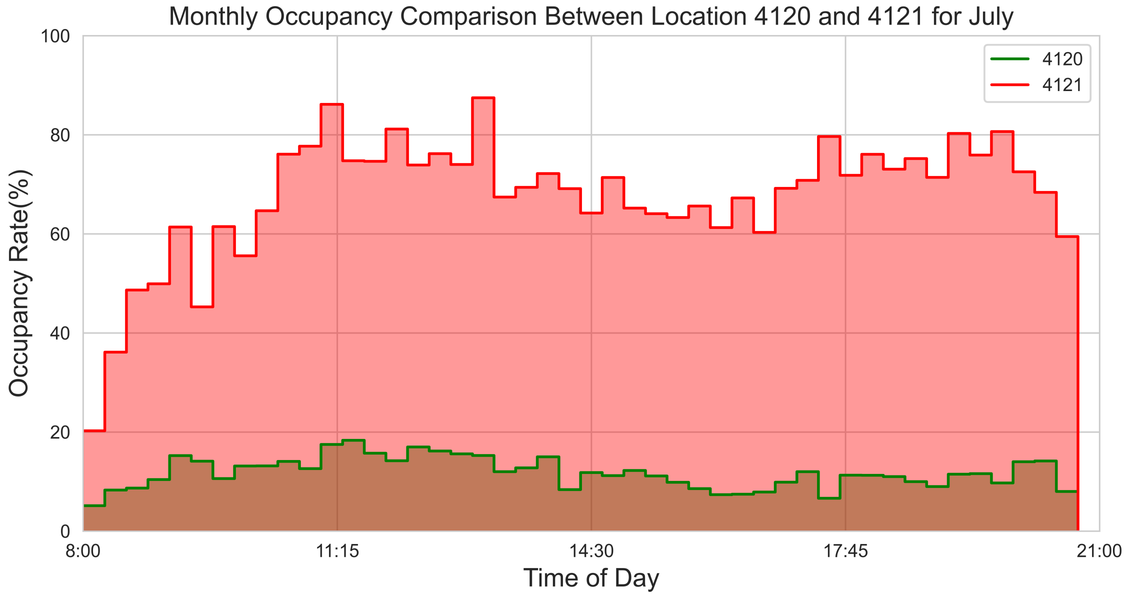

Figure 24 - Comparison between Two Parking Areas on Wellington St

PRICE ADJUSTMENT OPTIMIZATION

Algorithm Overview

This report developed an algorithm that systematically adjusts parking pricing by identifying parking zones where rate changes could balance the overall occupancy rates across the parking network. The algorithm selects a candidate zone for rate modification and uses a Machine Learning model (i.e., a Graph Neural Network - explained in the following) to predict the network-level outcome. Figure 25 outlines the algorithm's workflow, which starts by defining a network of parking zones. Figure 26 illustrates the parking network in Toronto, ON, where each parking zone is connected to its neighbours by a link if they are within walking distance (i.e., their Euclidean distance is less than 300 meters). The links depicted in Figure 26 allow the algorithm to extract spatial information such as pricing levels and the availability of nearby parking spaces. Some nodes in Figure 26 are not connected to any others; these are nodes where a potential price change might not be advisable, as there would be no alternative parking options for users. However, the algorithm is flexible and can override this if needed by adjusting the distance connectivity threshold.

The algorithm identifies candidate zones for pricing adjustments by conducting a statistical test (i.e., a T-test) at a 1% significance level across each link in the network. The T-test determines whether the mean occupancy rates between each pair of the connected zones are statistically different. If two zones have a statistically meaningful difference in mean occupancy rates, the link connecting them is marked red, while all other links are marked green. The algorithm uses the number of red links connected to a zone to predict the potential impact of a price change. The presence of more red links indicates a higher opportunity to reduce discrepancies in occupancy rates, suggesting that the zone may be a good candidate for a price adjustment. Each red link connects two zones: one with a lower occupancy rate and one with a higher occupancy rate. The algorithm is designed to increase prices only at zones with higher occupancy rates. The algorithm tests $0.50 increases at these zones, aiming to balance the occupancy rates across the network by shifting demand to less occupied zones. Prices are increased sequentially, and the impact is simulated using the machine learning model to predict changes in occupancy rates and link classifications. Once a price change is established by the algorithm, the neighbouring zones are also updated to reflect the impact of price changes. The algorithm continues until either the runtime limit is reached or no red links remain.

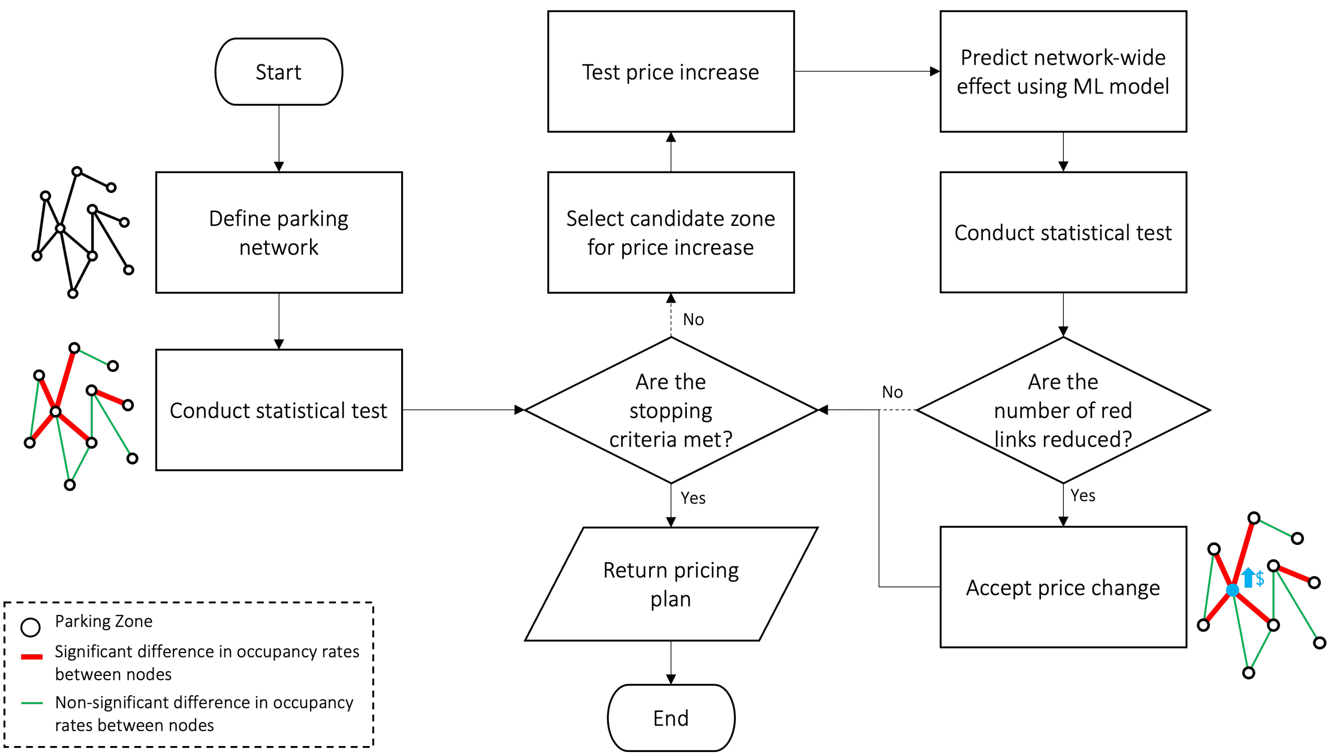

Figure 25 - Algorithm flowchart

Graph Neural Network

This report leverages a Graph Neural Network (GNN) to predict the outcome of each proposed price change. A GNN is an Artificial Neural Network designed to operate on graph data structures, enabling it to learn from the relationships and features associated with nodes (individual parking zones, e.g., zone capacity), edges (connections between zones, e.g., arc length), and the graph level (overall network structure, e.g., weather conditions). A GNN learns directly from a graph's structure and propagates features across the network. This is achieved through a process called *message passing*, where information is shared between connected nodes, allowing the GNN to capture the spatial and relational dependencies within the network. Hidden layers in the GNN process this information and apply non-linear transformations to learn complex patterns within the data. This helps the model to aggregate and propagate features across the graph (see Figure 27).

Figure 26 – Parking network in Toronto, Ontario (zoomed-in).

The proposed GNN in this report is trained to predict the number of open parking spaces, average vehicle parking duration, and number of vehicles arriving at each zone (i.e., demand for parking zone), using independent variables such as land use, proximity to the points of interest, weather conditions, and the number and pricing of adjacent parking spots. The three target variables together (i.e., the number of open parking spaces, average vehicle parking duration, and number of vehicles arriving at each zone) provide a comprehensive understanding of the system's response to price changes, as evaluating them individually would not capture the full impact. For instance, a price increase might leave demand unchanged but reduce parking duration, or vice versa.

The proposed GNN has a novel design where it parameterizes the features of a zone and its neighbours with different weights. Extensive testing of various GNN architectures revealed that the spatial location of a parking spot influences its demand. Certain spots, particularly those situated near the downtown core or other high-demand areas, exhibit stable occupancy rates despite price increases due to their strategic importance. By assigning different weights to the features of a zone and its neighbours, the model captures these spatial dependencies and interactions within the network, resulting in a more accurate prediction of occupancy patterns.

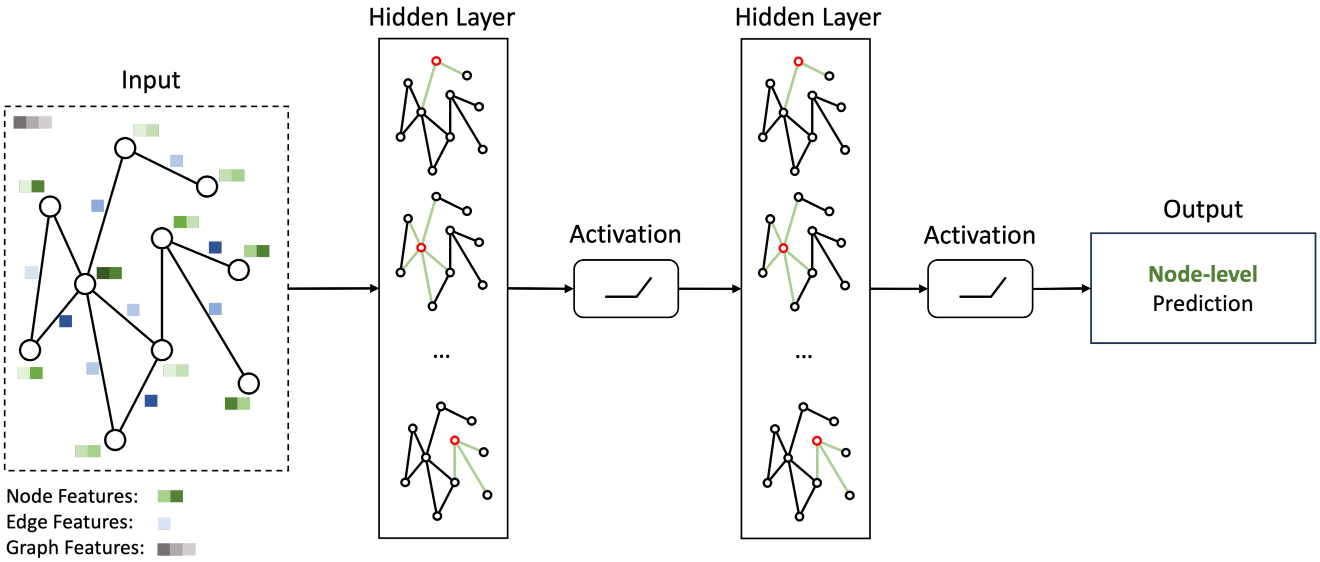

Figure 27 - The visualization of the proposed Graph Neural Network structure.

RESULTS

Graph Neural Network

To evaluate the performance of the GNN model, a holdout strategy was employed to split the data into training and testing sets. Specifically, the dataset was divided into an 80/20 ratio, ensuring that the splitting process was stratified based on the year of the data. The data includes records from both before and during the COVID-19 pandemic (i.e., 2019 and 2022), and the stratification process helps maintain a balanced representation of different years in both the training and testing sets. This is because statistical analysis showed that the COVID-19 pandemic impacted parking patterns and demand, hence, training and testing on balanced datasets would ensure that the GNN is more robust and has generalizable predictions.

For training the GNN, Mean Squared Error (MSE) was used as the loss function. MSE is a metric used for regression tasks that measures the average of the squared differences between the predicted values and the actual values. In the context of the GNN model, where the target variable comprises three components (i.e., parking demand, parking duration, and number of open spaces), the MSE loss function calculates the squared differences for each normalized component separately, summing these differences and then averaging the result over all the data points. Normalizing the components ensures that each component contributes equally to the loss calculation, preventing any single component from disproportionately influencing the model training due to differences in scale.

Table 5 assesses the prediction accuracy of the GNN and compares its performance on the training and testing data during each period of the day. The results indicate that the model performs consistently across both datasets, suggesting that it is not overfitted. Overfitting occurs when a model learns the training data too well, capturing noise and outliers, which results in poor generalization to new, unseen data. For instance, in the morning period, the training MSE is 1.508, and the testing MSE is 1.510, showing that the model has learned the patterns in the training data and generalizes well to the testing data. Similarly, in the afternoon and evening periods, the MSE values for training and testing sets are reasonably close.

Price Optimization Algorithm

Table 6 illustrates the results of the algorithm applied to the parking network. On average, with minor changes in the prices of the parking zones, the number of red links, which indicate zones with statistically significant differences in their occupancy rates, was reduced by approximately 2%. For all periods of the day, fewer than 40 zones underwent a price change. Most zones experienced a $0.50 hike, while a few had a $1.00 increase. The zones with the $1.00 increase are primarily located within the downtown core or near points of interest.

Table 5 – Graph Neural Network's performance during training and testing

Period	Mean Squared Error	
	Training Set	Testing Set
Morning (7 AM - 12 PM)	1.508	1.510
Afternoon (12 PM - 5 PM)	3.796	3.810
Evening (5 PM - 9 PM)	2.525	2.731

Table 6 - Number of zones with price changes and red links before and after price adjustments for different periods. The parking network has 2572 arcs in total (including green and red).

Period	Network	No. of zones with a price change	No. of red links
Morning (7 AM - 12 PM)	Before Price Change	-	2288
	After Price Change	35	2246
Afternoon (12 PM - 5 PM)	Before Price Change	-	2338
	After Price Change	36	2296
Evening (5 PM - 9 PM)	Before Price Change	-	2318
	After Price Change	38	2263

Table 7 provides detailed insights into the effects of increased parking prices on occupancy and demand across different times of the day. In the morning period (7 AM-12 PM), the results show an increase of 3.3% in the number of open parking spaces, resulting in an average of 9.1% decrease in the occupancy rate. This period is characterized by commuters heading to work, where late arrivals can result in penalties. Consequently, individuals are more inclined to pay for higher-priced but available parking spots to avoid tardiness. The increase in open spaces in the morning encourages people to park in these higher-priced zones, with a notable 40% shift towards these areas.

In the afternoon (12-5 PM), the number of open spaces increases by 8.7%, and the occupancy rate decreases by 12.9% in the target zones (i.e., zones with a price change). Additionally, both parking demand and the duration of parking usage see a reduction of 8.2%. The results show that the users seek more affordable parking spots during the afternoon. The evening period (5-9 PM) presents a different scenario from the previous periods, where the occupancy rate decreases by 6.7%, even though the number of open spaces remains almost unaffected. This period shows that the decrease in occupancy rate encourages people to stay longer, with a 40.8% increase in extended stays and a 44.7% rise in demand, often for activities like shopping or dining out. The unchanged number of open spaces, despite a lower occupancy rate, highlights the evening's lower baseline demand, where even a small absolute increase in demand results in significant percentage changes.

Table 7 - Summary of key metrics before and after price changes, averaged over each period for parking zones that underwent a price change.

Period	Network (Target Nodes)	Open Spaces	Occupancy Rate	Parking Duration (minutes)[1]	Parking Demand	Revenue (million CAD)[2]
Morning (7 AM - 12 PM)	Before Price Change	10.1	27.4%	81.6	10.5	31.72
	After Price Change	10.5	24.9%	80.6	14.7	31.75
	% Change[3]	+3.3%	-9.1%	-1.1%	+40.0%	+0.10%
Afternoon (12 PM - 5 PM)	Before Price Change	16.4	25.3%	72.9	37.0	20.60
	After Price Change	17.8	22.0%	67.0	33.9	20.62
	% Change	+8.7%	-12.9%	-8.2%	-8.2%	+0.11%
Evening (5 PM - 9 PM)	Before Price Change	17.3	7.4%	28.2	4.8	9.54
	After Price Change	17.3	6.9%	39.7	6.9	9.56
	% Change	+0.0%	-6.7%	+40.8%	+44.7%	+0.21%

[1] The average parking duration per vehicle during period-of-day.
[2] Annual revenue from parking zones across the entire network.
[3] The percentages indicate the relative change before and after the price adjustments.

INTERACTIVE WEB-BASED TOOL FOR PARKING OPTIMIZATION

Key Features of the Tool

1. Visualization: Users can display various parking metrics on the map, such as occupancy, capacity, occupancy ratio (occupancy divided by zone capacity), number of open spaces, and hourly price. Additionally, the map showcases different wards, highlighting the impact of varied land uses and demographics on parking patterns. The users can visualize metrics aggregated over space (i.e., over ward zones or single parking area) and time (i.e., morning, afternoon, evening or as an average across these periods).
2. Interactive Manipulation: Users can interactively select any parking location and modify attributes such as price, capacity, and maximum allowable parking duration. The tool then simulates the impacts of these changes on the surrounding zones using the algorithm presented.
3. Auto Optimization: The tool provides an 'optimize' feature, where users can specify the number of parking zones (or zones per ward) for price modification and the allowable range for adjustments (see Figure 29). Upon activation, the tool autonomously calculates and implements optimal pricing adjustments across the selected areas.

In addition to the interactive tool, the report also developed a dashboard using ArcGIS that offers increased flexibility and enhanced data manipulation capabilities (as seen in Figure 30). This dashboard serves as a supplementary visualization platform. Users can plot and observe trends directly on the map, providing a snapshot of the current state without the optimization features offered by the main tool.

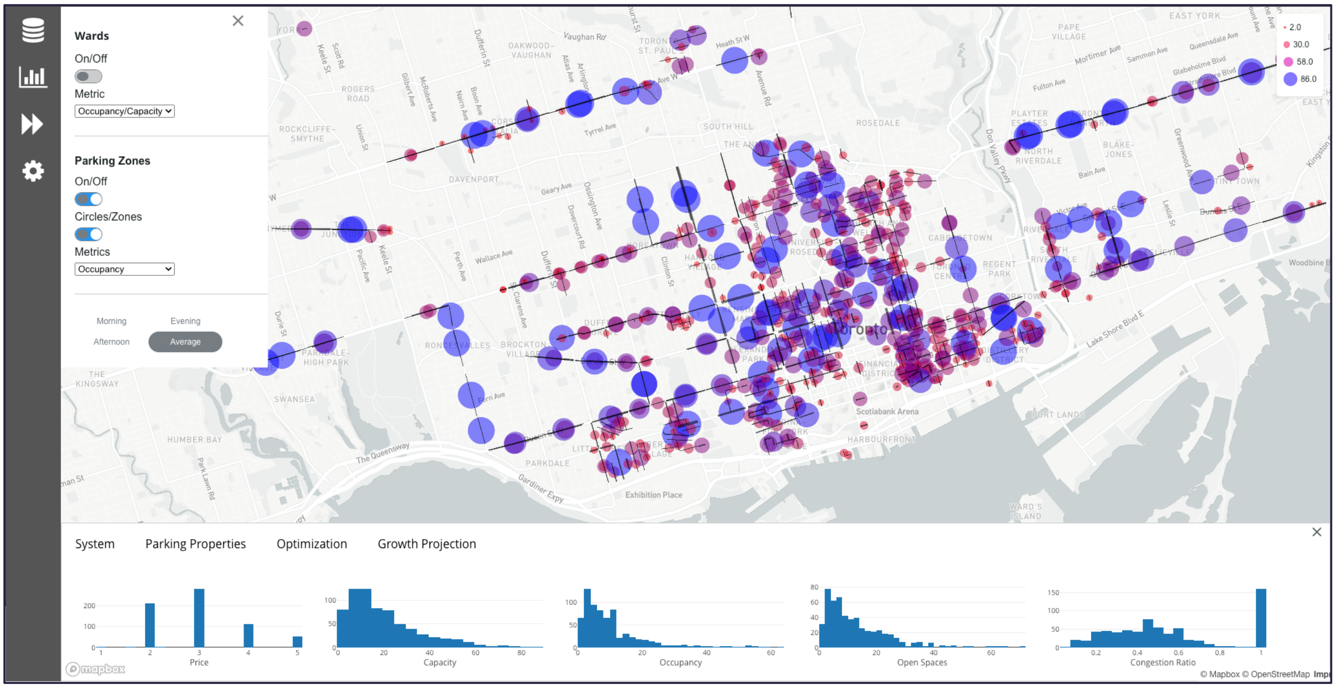

Figure 28 – Interactive parking management tool

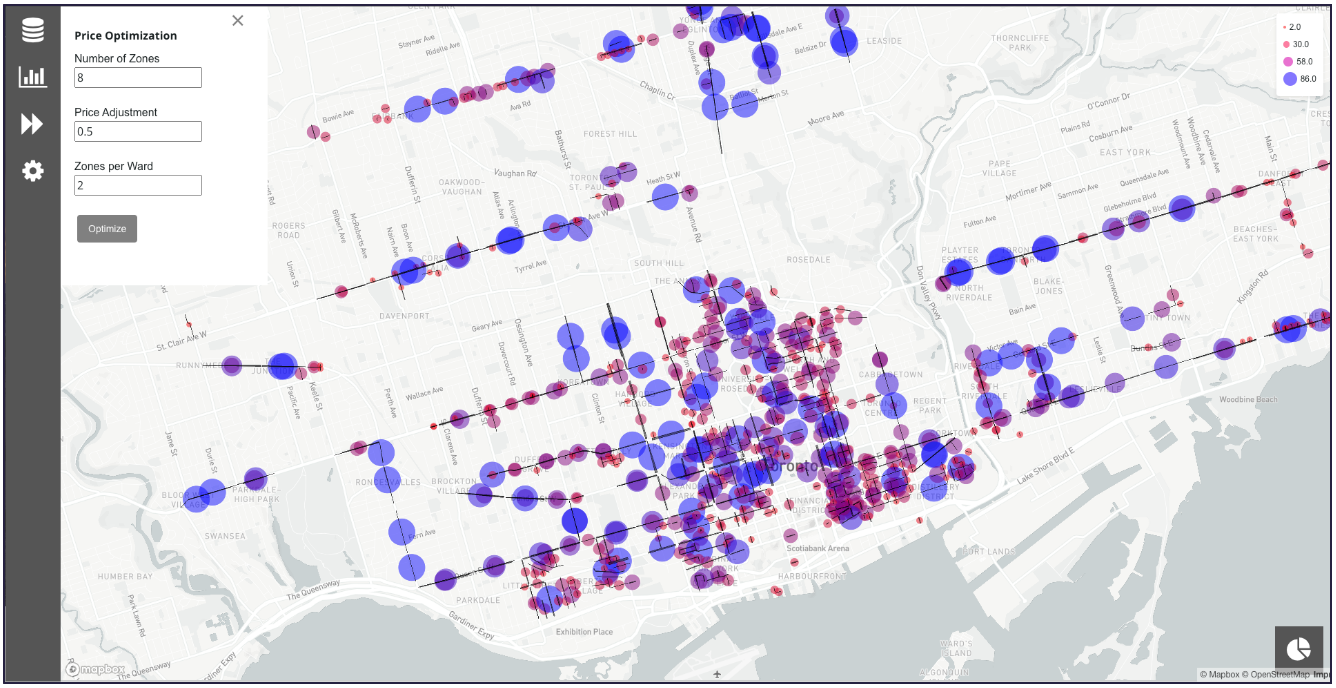

Figure 29 – Auto price optimization feature on the interactive parking management tool

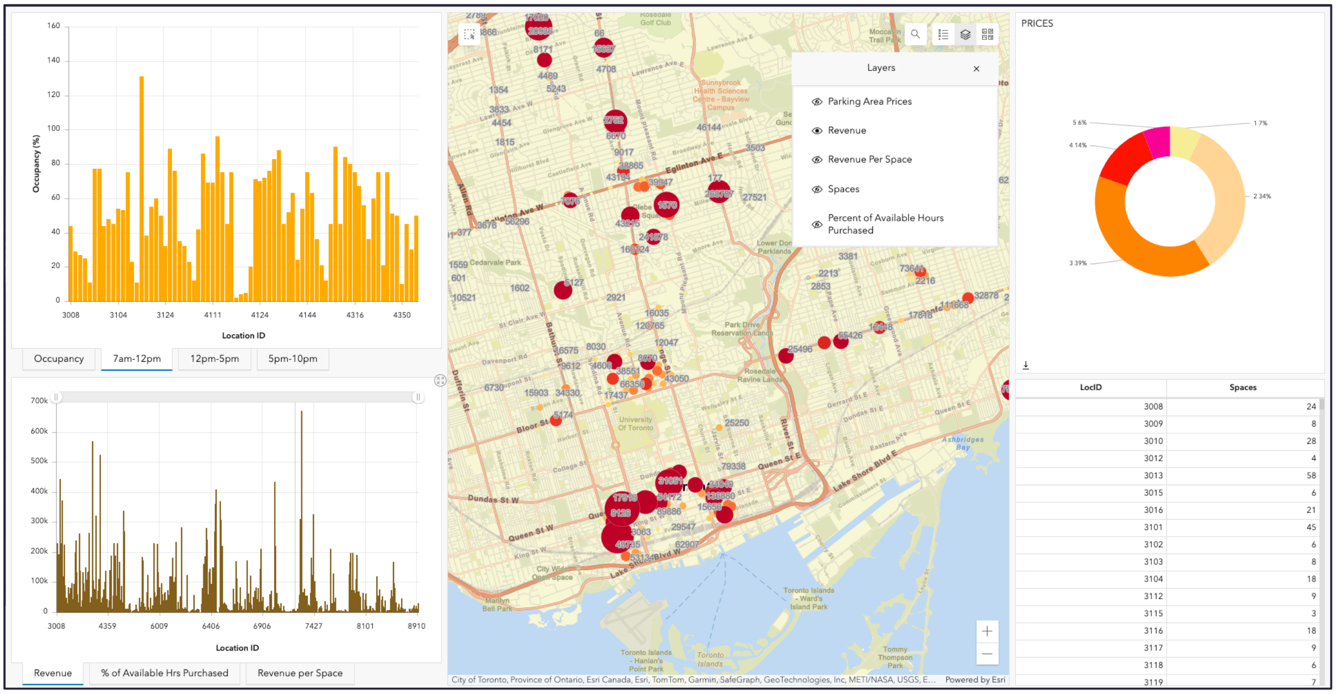

Figure 30 – ArcGIS parking management dashboard

CONCLUSIONS

Managing on-street parking in densely populated urban areas is a multifaceted challenge due to the high demand for limited parking spaces. With parking zones becoming full, this prolongs the search time for drivers, adding congestion on the road. Implementing a dynamic parking pricing policy is an innovative approach to address this complexity. Dynamic parking pricing levels the occupancy of parking zones by adjusting parking prices according to demand to ensure parking spaces are open in popular parking zones while also attracting drivers to underutilized locations. Cities like San Francisco, Washington DC, and Los Angeles have seen success in dynamic parking pricing pilots, improving parking availability, reducing parking search times, and decreasing congestion. From the positive effects of the pilots, cities are now expanding the parking pricing policy within their cities. However, scaling dynamic parking pricing is expensive due to the cost of occupancy detection infrastructure. Cities often deploy sensor technologies to gather occupancy rates for price adjustments. Maintenance of these sensors becomes a concern, as some sensors see early battery failure, and electromagnetic interference reduces occupancy rate accuracy. Using transaction data to predict parking occupancy becomes a viable solution to the dynamic occupancy detection infrastructure problem, as this data is widely available and manipulable. Installing occupancy detection infrastructure can be avoided, as a historical transaction data prediction model can generate future parking occupancies based on price changes.

Preliminary analysis was conducted on the transaction data to gain insight into trends and patterns. The analysis revealed that the day of the week impacted demands, as drivers tended to park at distinct locations depending on whether it was a weekday or weekend. Weekdays exhibited a surge in morning occupancy, due to commuters parking before work. Conversely, weekends showed higher afternoon occupancy rates, linked to leisure activities. Evening occupancy rates remained consistently low throughout the week, suggesting reduced demand for on-street parking during these times. Furthermore, the analysis revealed that price hikes in certain parking zones, particularly those near points of interest or the downtown core, do not significantly alter demand or parking duration. The developed model reflects these spatial variations in demand reactions.

Overall, the proposed pricing strategy involves a dynamic system that adjusts prices based on parking network occupancy. This aims to optimize spaces in high-density areas while filling low-density ones, predicting nearby changes due to price adjustments. The pricing adjustments demonstrated improvements in parking availability and balanced occupancy rates. Specifically, in the morning period (7 AM – 12 PM), price adjustments led to a 3.3% increase in open parking spaces and a 9.1% decrease in the average occupancy rate, despite a 40.0% increase in parking demand, attributed to vehicles parking for 1.1% shorter durations. In the afternoon period (12 PM – 5 PM), open spaces saw an increase of 8.7%, with occupancy rates reduced by 12.9% due to the 8.2% reduction in parking demand and duration. During the evening period (5 PM – 9 PM), open spaces increased slightly by 0.008%, resulting in a 6.7% reduction in occupancy rates, while parking demand and duration surged by 44.7% and 40.8%, respectively. Minor price adjustments, typically $0.5 increments, and occasional $1.00 raises in high-demand areas, were employed. Notably, areas closer to points of interest justified higher prices due to their spatial demand bias.

BIBLIOGRAPHY

City of Baltimore. (2017, November 9). *Demand-Based Parking Meter Rate Setting*. Parking Authority. https://parking.baltimorecity.gov/demand-based-parking-meter-rate-setting

City of Boston. (2016, November 9). *Performance Parking pilot*. Boston.Gov. https://www.boston.gov/sites/default/files/embed/p/performance_parking_final_report_-_web_1.pdf

City of Calgary. (2022). *Calgary On-Street Parking Rates*. https://www.calgaryparking.com/findparking/onstreetrates

City of Seattle. (2017). *Paid Parking Rates—Transportation | seattle.gov*. https://www.seattle.gov/transportation/projects-and-programs/programs/parking-program/paid-parking-information/street-parking-rates

LA Express Park. (2012, May 21). *LA Express Park*. https://content.cdntwrk.com/pdf_viewer/web/viewer.html?file=https://content.cdntwrk.com/files/aT0xNDMyMTg3JnY9MSZpc3N1ZU5hbWU9bGEtZXhwcmVzcy1wYXJrLXBhcmtpbmctc21hcnRlci13aXRoLWRlbWFuZC1wcmljaW5nJmNtZD1kJnNpZz02MmJhOGVhNTc4OThkYmQ4YzkxNjM0N2M0N2EzZjZiZQ%25253D%25253D

ParkDC. (2019). *ParkDC*. https://rosap.ntl.bts.gov/view/dot/56382/dot_56382_DS1.pdf

SFPark. (2013, January 25). *SFpark Pilot Program* [Text]. SFMTA; San Francisco Municipal Transportation Agency. https://www.sfmta.com/sites/default/files/reports-and-documents/2018/08/sfpark_pilot_overview.pdf