VU Research Portal

Assessing the predictive value of traffic count data in the imputation of on-street parking occupancy in Amsterdam

Calvo, Pablo Martín; Schotten, Bas; Dugundji, Elenna R.

published in

Transportation Research Record
2021

DOI (link to publisher)

[10.1177/03611981211029644](https://doi.org/10.1177/03611981211029644)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Calvo, P. M., Schotten, B., & Dugundji, E. R. (2021). Assessing the predictive value of traffic count data in the imputation of on-street parking occupancy in Amsterdam. *Transportation Research Record*, 2675(12), 330-341. <https://doi.org/10.1177/03611981211029644>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Assessing the Predictive Value of Traffic Count Data in the Imputation of On-Street Parking Occupancy in Amsterdam

Transportation Research Record
2021, Vol. 2675(12) 330–341
© National Academy of Sciences:
Transportation Research Board 2021
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/03611981211029644
journals.sagepub.com/home/trr

Pablo Martín Calvo¹ , Bas Schotten² , and Elenna R. Dugundji^{3,4}

Abstract

On-street parking policies have a huge impact on the social welfare of citizens. Accurate parking occupancy data across time and space is required to properly set such policies. Different imputation and forecasting models are required to obtain this data in cities that use probe vehicle measurements, such as Amsterdam. In this paper, the usage of traffic data as an explanatory variable is assessed as a potential improvement to existing parking occupancy prediction models. Traffic counts were obtained from 164 traffic cameras throughout the city. Existing models for predicting parking occupancy were reproduced in experiments with and without traffic data, and their performance was compared. Results indicated that (i) traffic data are indeed a useful predictor and improves performance of existing models; (ii) performance does not improve linearly with an increase in the number of counting points; and (iii) placement of the cameras does not have a significant impact on performance.

On-street parking can be found in most cities around the world. It is a key service for urban mobility, whereby drivers can conveniently reach many locations of the city with their vehicles door to door. Usually being a public service, authorities are responsible for executing appropriate policy making to optimize this scarce resource. One of the main concerns is the increase in traffic congestion owing to high occupancy of on-street parking space (1, 2). This relationship serves as the seed for the motivation of this research: if high parking occupancy leads to traffic congestion, could traffic data serve as a predicting feature for parking occupancy?

This paper proposes assessing traffic data as a predictive feature for parking occupancy. This is observed in two different cases: in imputation models for completing the sparse spatial-temporal view of the city obtained through probe vehicles and in forecasting the future occupancy in short-term time steps. Imputation models are relevant for cities that have chosen probe vehicles as their occupancy monitoring tool, a promising solution to the parking occupancy monitoring problem in relation to effectiveness and cost (3–5). This is because the sparsity of the monitoring leads to data gaps across time and space. Without imputing those gaps, authorities are left with a partial view of occupancy. Short-term forecasting

models are required when authorities want to decide on policies for the near future (6).

This leads to the following research questions:

- Can sensor-gathered data about traffic flows reduce prediction error in existing imputation and forecasting models for on-street parking occupancy in the city of Amsterdam?
- How does the number of sensors affect the prediction error?
- How does the geographical position of sensors affect the prediction error?

The city of Amsterdam, in the Netherlands, was chosen for experiments given the access to both data and models related to the topic provided by the municipality. To answer empirically the proposed research questions, traffic counts obtained through street cameras were

¹University of Amsterdam, Amsterdam, The Netherlands

²City of Amsterdam, The Netherlands

³Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

⁴Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

Corresponding Author:

Pablo Martín Calvo, pabломartincalvo@gmail.com

acquired, and the predictive model proposed by Schmidt was reproduced with and without using this data as a predicting feature (7). The results of these experiments showed the impact of the traffic data on the performance of the model. Additional experiments provided insight into the impact of camera quantity and placement.

The rest of the paper is organized as follows: Related Work provides an overview of the existing research on the need for occupancy data in parking economics and a selection of data and methods that have been tried presented in literature. Methodology presents the data, model, and evaluation approach that were used in this research. Results presents the empirical results obtained through the performed experiments. The Discussion section contains reflections on the knowledge that can be drawn from the obtained results and their implications. The Conclusion section presents the conclusions of the paper.

Related Work

Related work is reviewed from three angles in the following subsections: the models and theories around on-street parking and how occupancy prediction is relevant, the relationship between parking occupancy and traffic congestion, and finally how models and data have been used in the past, with a special focus on the works that have attempted to use traffic data.

On-Street Parking Economics and Policy

On-street parking, also known as curb parking, has been modeled many times with some formulation of total social welfare as the variable to maximize through parking policy (8–10). This approach allows for taking into account all of the economical aspects of parking as well as its main externalities, such as space occupation, increased traffic, or time wasted cruising. (Cruising refers to the situation in which a driver cannot find a free spot in his destination and keeps driving around the area while waiting for a spot to become available.) Some authors have given explicit guidelines related to occupancy: Shoup recommends aiming for having all blocks almost full, but not full. This leads to a drastic reduction in cruising, which in turn improves the resulting overall welfare (1).

Pricing is often proposed as one of the main tools for pursuing target occupancy levels. Several authors agree that on-street parking is underpriced in most cities in the world (1, 8, 9) and that prices should be increased to reach the aforementioned occupancy goals. When differences in time and space for parking demand are also taken into account, the need for dynamic spatial-temporal pricing appears. This idea was already proposed decades ago by Vickrey, but technological and

data collection constraints made this idea infeasible at the time (11). However, nowadays the state of technology allows for applying Vickrey's proposals. In 2011, in San Francisco, the SFPark project deployed a pricing system that adjusted parking meter prices once every 6 weeks depending on the mean occupancy being in a specific range (60% to 80%) (6). Different prices were set across blocks, different times of the day, and weekdays and weekends. After a year of activity, the average price across the entire project area was the same, but the spatial patterns allowed for 67% of the previously underused areas to increase occupancy, whereas 68% of overoccupied blocks reduced their average occupancy level.

The main takeaway from this section is that there is plenty of work to be done around on-street parking policy and the social benefit to be obtained through it, but this is only possible with data that is both accurate and complete across time and space.

On-Street Parking and Traffic

On-street parking and traffic have a close relationship. Several studies have shown that significant proportions of traffic congestion can be produced by drivers cruising while searching for a spot. Shoup listed a total of 16 studies on cruising performed during the 20th century (1). The aggregated results of these studies show that, in the worst conditions, drivers take on average 8.1 min to find a suitable spot and 30% of the traffic volume is cruising for a spot. Van Ommeren et al. worked on measuring the same problem for the Netherlands (12). In their 2011 paper, empirical data from the Dutch National Travel Survey on cruising behavior for the Netherlands was studied. Their analysis of the data highlighted that 30% of car trips end with some time spent cruising, with a mean cruising time of 37 s. This low value could be explained by the way that many Dutch local governments enforce on-street prices as high as the off-street garages, relatively discouraging on-street parking. The authors also noted that the mean cruising time is 3 s longer in cities of above 100,000 inhabitants. Unfortunately, no specific details were given for downtown areas of the major Dutch cities, where cruising is often assumed to be a much bigger problem than in other areas.

All this knowledge serves as motivation for the main hypothesis of this paper: that information about parking occupancy can be gained from traffic data, and thus it could be a useful predictive feature.

Imputation and Forecasting Methods and Traffic as an Explanatory Variable

In this section, two distinct problems are reviewed: imputation of parking occupancy in areas and times that have

not been monitored (as in the Amsterdam case) and forecasting of future parking occupancy. Previous attempts at using traffic data for these challenges are also reviewed.

Imputation Methods. Imputation methods for the case of sparse occupancy monitoring are rare, as Schmidt pointed out in his recent work on the topic (7). A few papers that deal with this issue directly, or with similar problems, have been identified.

Bock is one of the authors that has worked on the topic. In his work with Di Martino and Sester in 2016, he evaluated the effectiveness of a fleet of probe vehicles by composing a synthetic occupancy dataset from the full sensor data of SFPark, assuming only a 30% observation rate across time and space (4). The work used a random forest classifier that could predict whether road segments would be full or not full in several future horizons, up to 60 min. This model used data from the same road segment as well as the neighboring ones. The authors concluded that, although the scenario with sensor monitoring (leading to a 100% coverage on space and time) had the best performance in the prediction results, the simulated probe vehicle data led to a very close performance. This seems to indicate that it is possible to execute forecasting with a sparse view of the occupancy across the city, and that predictions for a specific road segment can be made with data from neighboring ones. The conclusion that nearby road segments can provide information on the occupancy of a specific one was also supported in another work by Bock and Sester in the same year (13).

Ionita et al. developed a model for predicting occupancy in completely unmonitored areas (14). The chosen approach was to cluster areas in the city based on what the authors called “city data,” which was composed with geolocated amenities and points of interest obtained from OpenStreetMaps. Once these clusters had been made, occupancy data from areas that were monitored, together with the developed clusters, were used to train a model that could output predictions for the areas that were not monitored. Different machine learning models were used, with experiments being executed using occupancy data from the SFPark project.

Forecasting Methods. Whereas the works in the imputation field for street parking are rather scarce, there are plenty of methods in the forecasting area. Performing a complete review of this field and all the existing literature would be beyond the scope of this paper. Thus, a selection of representative works will be presented.

Pullola et al. also tried to predict the occupancy of a parking lot, as part of a wider GPS system for recommending optimal parking destinations for drivers (15). The future occupancy of the parking lot is modeled as a

Poisson process. No evaluation was performed on real data, with only a synthetic dataset being used for showcasing the behavior of the proposed system.

Ji et al. developed a wavelet neural network for predicting the available parking spaces in parking lots in the short term (16). The model is designed to use previous occupancy data from the same parking lots. An evaluation was performed using data from a parking lot in Newcastle, showing promising results with little error, measured as Mean Squared Error.

Zheng et al. compared the result of several machine learning models on a dataset from the SFPark project and a dataset from the city of Melbourne (17). Their goal was to predict the occupancy rate of street parking in 15-min steps. The compared models were regression trees, support vector regression, and neural networks. Using only previous occupation data, they concluded that regression trees had the least error out of the three compared models.

Rajabioun and Ioannou proposed a vector autoregressive model intended for both on-street and off-street parking, which accounted for both temporal and spatial relationships between the different parking areas and points in time (18). The forecasts were produced for short-term periods under an hour. The model was tested with data from the SFPark project for validation.

Badii et al. developed several forecasting models (BRANN, RNN, ARIMA, and SVM) for predicting available spaces in gated garages and tested them on several garages of the city of Florence, Italy (19). A mixture of data sources, obtained from the city open-data system, were used as predictive features for the models.

Finally, Fan et al. recently proposed a long short-term memory (LSTM) model based solely on time-series data, which could execute multistep, short-term forecasts on garage vacant space (20). The model was benchmarked against other common machine learning models with real time-series data obtained from two garages. The obtained results showed extremely accurate forecasts for short time spans, with a Mean Absolute Error value below 5% for forecasts with a horizon of 15 min.

Use of Traffic Data. Finally, three papers in which traffic data were considered as a potential feature for predicting parking occupancy are presented. These deserve a special focus since they are highly related to this paper. All three focused on the forecasting problem. No research was found in which traffic data were used for imputation purposes.

The earliest work in which traffic data were used as a predicting variable for parking occupancy was developed by Ziat et al. in 2016 (21). In their paper, the authors proposed a model to solve the problem of “cross-forecasting of road-traffic prediction and parking occupancy.” Their

motivation was similar to the authors as they stated that “a broader approach encompassing both parking and traffic could lead to better prediction because of the interrelation between traffic fluidity and parking availability.” The main difference with the current research is that whereas here parking occupancy is the only concern, their proposed model made forecasts for both traffic congestion and parking occupancy. The presented model was a multilayer perceptron, and it was tested with sensor-obtained data from the 50 busiest roads and the 30 busiest garages in the French city of Lyon. The empirical results of their experiments showed that their proposed model’s forecasts had a smaller error than other baselines. And the main finding that is relevant to our work is that the error on parking occupancy prediction was significantly lower when using both parking and traffic data, as compared to simply using parking data. This seems to confirm that traffic data can indeed improve predictions on parking occupancy. An important aspect to keep in mind is that the authors worked with garages, and not on-street parking, so whether the same effect could be obtained for street parking was uncertain.

A second work in 2019 by Yang et al. also attempted to use traffic data (22). In this study, the authors’ goal was to develop a model for forecasting street parking occupancy on a block level with a 30 min horizon. The proposed model is a complex deep learning design that includes usage of graph convolutional- (GCNN), LSTM- and multilayer feed forward neural networks. One of the innovative proposals of the authors was to model the spatial dependencies of data through the use of graphs and GCNN. The final output of the model is the forecasted occupancy for every block of parking spaces. The model is designed to be able to flexibly include multiple data sources, and in the scope of the paper several are used: parking meter transactions, road network, weather, parking occupancy, and traffic speed. The model was evaluated with data from the downtown area of the city of Pittsburgh. The general results of the experiments were promising, since all proposed baselines were outperformed by a significant margin. Focusing on the effect of traffic speed, the experiments showed a significant improvement when using traffic congestion data in comparison to not using it. The authors valued the effects positively and concluded that “The speed profile of the network reflects the real-time demand and congestion of both parking and traffic flow. This confirms our assumption that the correlation between traffic congestion and parking occupancy is significant.”

Finally, a 2012 paper by Hössinger et al. proposed a model that aimed to predict occupancy for short-term time horizons up to 90 min ahead (23). Initially, the authors considered traffic count data as a predicting feature, but this was discarded during exploratory analysis.

They stated that “although we found some large correlations among the numerous candidates, they have different signs and don’t reveal any recognizable geographical pattern.” Because of this, they deemed the relationship between traffic and parking occupancy “too fragile” and decided not to include it in their final model. This statement seems to be in conflict with the views and findings by Ziat et al. (21) and Yang et al. (22). It is worth observing the following statement from the paper by the latter that managed to successfully use similar data: “Clearly, it is inappropriate to model the relationship between traffic speed and parking arrivals via models with low complexities, while a deep neural network may have the potential to learn this complicated relationship.” The outcomes of both works seem to indicate that traffic and parking occupancy have highly complex relationships across time and space that might not be easy to spot through conventional methods.

The works by Ziat et al. (21) and Yang et al. (22) show successful use of traffic data, both traffic counts and speeds, in forecasting future occupancy. Can these results be translated to Amsterdam’s case successfully? And could a similar performance be achieved in imputing values for the sparse spatial-temporal view that the City of Amsterdam obtains through its probe vehicles?

Methodology

Different combinations of data sources, preprocessing techniques, and models were brought together in several experiments to obtain empirical measurements on the effectiveness of traffic data for estimating parking occupancy. Two baselines were proposed to compare performance with the obtained results. The following subsections describe these components in detail.

Data Sources and Preprocessing

Parking Occupancy. In this paper, parking occupancy (also known as parking pressure, or simply occupancy for the sake of brevity) is defined as the fraction of on-street parking spots that are occupied by vehicles as part of the total number of available parking spots. This needs to be bound to a certain geographical area, a , at a point in time, t .

$$occupancy_{t,a} = \frac{occupied_{t,a}}{available_{t,a}} \quad (1)$$

In this context, “available” means that the spot can be used by a vehicle, regardless of whether it is actually being used. In contrast, parking spots are unavailable when the municipality of Amsterdam temporarily forbids parking in them, such as when a house moving

company requests a temporary permit to use those spots to park their trucks.

Occupancy was the target data in the executed experiments: in all of them, occupancy at a certain point in space and time was either imputed or forecasted. Occupancy observations in this work were obtained by grouping several “point occupancies.” A point occupancy (p) is an occupancy value for a specific road segment, r , at a certain point in time, t . Point occupancies were obtained through a combination of scan observations provided by Amsterdam’s probe vehicles, maps of the existing parking spots in the city maintained by the municipality, and geographical descriptions of the road network obtained through OpenStreetMaps.

The chosen geographical space unit for the experiments was the neighborhood. Neighborhoods are one of Amsterdam’s official administrative spatial divisions (24). The main motivation for using this unit was to maintain comparability with the research done by Schmidt (7). Furthermore, several benefits make this a suitable unit: it covers the entire city of Amsterdam seamlessly, it is an official and widely adopted division of space, and it strikes a balance between areas that are too expansive, like the entire city, which would make results uninteresting, and areas that are too small, such as individual road segments, which would make results unstable and extremely sparse in relation to space. Most neighborhoods have areas between 0.01 and 2.5 km², with a few outliers being larger than 2.5 km².

In the time dimension, data were obtained by grouping point occupancies in hourly intervals. Specifically, occupancy was computed as the mean of the set of point occupancies, P , observed in a road segment, r , belonging to neighborhood, b , and at time, t , contained in the span of hour, h . Since point occupancies were observed for road segments of different sizes in relation to total parking spots, for fairness, the mean was weighted by the number of spots available in each road segment. The resulting values indicated the mean occupancy of every neighborhood during each hour interval,

$$\begin{aligned} \text{occupancy}_{b,h}(P) &= \frac{\sum_p^P \text{occupied}_p}{\sum_p^P \text{available}_p} \\ \text{s.t.} \\ r_p &\in b, \forall P \\ t_p &\in [h, h + 1), \forall P \\ p &\in P \end{aligned} \quad (2)$$

Only parking spots belonging to the *fiscaal* category were selected. Fiscaal spots require a valid parking right, be it long-term or a temporary ticket, and thus are the only ones that are monitored by the city’s probe vehicles. This is because the original purpose of this probe vehicle fleet

was to enforce parking regulations. The generation of the occupancy data used in this project was an unintended result of their main activity.

The dataset for this paper was provided by the parking department of the City of Amsterdam. These data were collected through the probe vehicles used by the municipality, which drive through Amsterdam monitoring paid parking. The exact schedules and policies of their routes are confidential and the City of Amsterdam requires them to be kept undisclosed. The dataset was composed by aggregating 4,262,693 point occupancies along 340 neighborhoods in Amsterdam. These were observed between 1 June 2019 and 28 February 2020, with 7 days lacking in between owing to missing data in the data sources. Some 27,934 were removed as a result of containing incomplete data of the scanned road segments; 49,003 observations were removed because they were labeled as outliers for having occupancy values higher than 120%. Occupancy values above 120% can occur either because of illegal parking on unauthorized spaces or because of several, smaller than usual vehicles being parked in the same street segment. The specific choice of 120% as the threshold was made to ensure consistency with Schmidt (7). The resulting aggregation yielded 323,377 occupancy observations per neighborhood and hour. This number decreased to 91,596 after removing neighborhood and hour combinations in which less than 80% of the neighborhood parking spots had been observed by the probe vehicles. The motivation for this was to ensure that the observations accurately estimated the actual occupancy of the neighborhood. The specific choice of an 80% threshold was also made to keep results consistent with the work by Schmidt (7). A final filter was executed to remove the combinations of neighborhood and hour that had fewer than 10 point occupancy samples. The remaining data comprise 61,269 observations. The sparsity of the data across time and space was very high: out of all the possible combinations of neighborhood and hour for the selected time span and neighborhoods, only 3.24% of them had observed values. The final data comprises, for each observation, a neighborhood code, a date and hour, and the corresponding observed occupancy.

Parking Tickets. Data from the Dutch National Parking Database were provided by the City of Amsterdam and used in this work to reproduce the predictive features proposed by Schmidt (7) with the aim of comparing their performance. The original motivation to use this data was the hypothesis that patterns in the acquisition of tickets could hold signals related to the state of parking occupancy in the city. (In the context of this paper, parking tickets, or simply tickets, always refers to temporary parking rights purchased at parking meters or through

Table 1. National Parking Database Feature Description, Where $h \in \{0, 1, 2, 3\}$

Columns	Description
Neighborhood code	Unique ID for the neighborhood
Date	Corresponding date for the data point
Hour (h)	Corresponding hour for the data point
Ticket count at time step $h-i$	Number of temporary parking tickets purchased per neighborhood (e.g., via parking machine or mobile app)
Ticket percentage at time step $h-i$	Number of purchased tickets divided by the total number of available parking spots per neighborhood
Ticket occupancy at time step $h-i$	Number of parking spots occupied by vehicles with purchased tickets, weighted by their duration as a proportion of the hour

mobile applications, and never to fines resulting from illegal parking.)

The features include several metrics related to the number of tickets sold and available in each neighborhood in the hours previous to the one that is to be predicted. Only some of the features were necessary for the experiments in our research. Preprocessing of the data and the building of these features were executed in the same fashion as Schmidt and resulted in the same features being available. Table 1 indicates the features that were used.

Traffic Counts. The aim of this study was to find out whether traffic data could be used as an explanatory variable for parking occupancy. Taking another look at the challenge, “traffic data” might come across as a rather vague term. What data are going to be used, out of the complex underlying reality of thousands of citizens driving around the city every day?

The two works by Ziat et al. (21) and Yang et al. (22) used two different types of traffic data: traffic counts in the former and mean speed of traffic in the latter. In this work, traffic counts were selected as the experimental data, given their accessibility from the City of Amsterdam (see Table 2). Other types of data, such as speed data or vehicle typology, could also be interesting features to explore are not in the scope of this paper.

For the experiments, a dataset with traffic counts throughout Amsterdam was obtained. This dataset was gathered through 336 cameras, known as mobility cameras, distributed throughout the city (see Figure 1). These cameras are used for several purposes, but one of the results is a database that records a new entry every time a vehicle drives through its line of sight. These counts were aggregated to hourly intervals, summing the total number of vehicles that were detected by the camera. Thus, the data span the time period from 1 June 2019 to 28 February 2020 and hold the number of vehicles counted by every camera each hour. Hourly bins were used to keep the traffic data at the same aggregation level as

Figure 1. Positioning of mobility cameras. Colored markers correspond to the four groups of cameras used in the experiment with geographical groups.

parking occupancy data. Shorter time steps were not explored but could be an interesting topic area for future research. No spatial aggregation was performed, keeping individual time-series for each camera. Out of the 364 cameras, 164 were selected because of having continuous data through more than 95% of the time span. The data from the other cameras were discarded.

From this main dataset, the following six different views were created to perform different experiments:

1. One hour: for every hour in the period, the corresponding counts from every camera.
2. Twelve surrounding hours: for every hour in the period, the counts of the previous 6 h and the next 6 h were appended. Thus, the feature set becomes the traffic count of the 12 h surrounding the target time step for the 164 cameras. The

motivation for this longer time window was the hypothesis that models could find more complex patterns in longer representations of traffic behavior.

3. Twelve previous hours: same as the 12 surrounding hours view, but using the 12 h preceding the time step instead of 6 h before and the 6 h after.
4. Random 50 cameras: same as the 12 surrounding hours view, but using 50 selected cameras instead of the 164 cameras. Five groups were built in this pattern. In all groups, the cameras were chosen at random, with equal probability for each candidate. Reported error metrics on this experiment correspond to the mean across these five random groups. This experiment was aimed at understanding the effects of low availability of traffic cameras.
5. Random 10 cameras: same as the 50 random cameras view, but using only groups of 10 cameras. The same selection procedure was used. Reported error metrics on this experiment also correspond to the mean across these five random groups.
6. Geographical groups: for this view, four groups of cameras were selected depending on their position in the city: north, south, east, and west (see Figure 1). Each group contained 10 cameras (except for the north, which had nine, owing to no more cameras being available in that area). Cameras within each area were selected randomly.

Feature extraction was performed on the raw traffic counts before being used for model training. The motivation for this was twofold: on the one hand, it was deemed possible that the high dimensionality of the data could lead to negative side effects related to the “curse of dimensionality” (25). On the other hand, reducing the dimensionality of data to a constant number of features made it possible to keep the structure of the model unchanged across experiments that used different views of the traffic data. The tool of choice to extract the features was a feedforward autoencoder.

Model

To maintain comparability with the findings by Schmidt (7), gradient boosting machines (GBM), one of the models used in that work, was chosen.

GBM was proposed by Friedman (26). This machine learning technique is suitable for both regression (desired for the task at hand) and classification. It belongs to the family of boosting algorithms (27). The algorithm sequentially generates weak learners, updating the weights of training observations in the loss function to increase those with the biggest error, and thus increases

focus on the worst performing areas of the target space. The implementation found in the scikit-learn package (28) was used in the experiments.

GBM was specifically chosen out of the several models evaluated in that work as it was the one that obtained the best performance among the candidates. Furthermore, as it was found to be best performing in the previous work (7), model training was performed on a neighborhood level. This means one independent model was created, trained, and evaluated for each of the neighborhoods in the data set, and was only trained with occupancy observations belonging to that neighborhood.

Finally, a hyperparameter choice was made and kept constant throughout the project. No tuning was considered necessary given that the task at hand was assessing traffic data, and not aiming at maximum performance in the predictions.

Performance Metrics and Evaluation

The task at hand was a regression problem. For every combination of neighborhood and hour, the percentage of occupied spots out of the total needed to be generated. This translated to the range between 0 and 1, with values superior to 1 being possible given that, occasionally, more vehicles than spots are found on the streets.

Given this, the chosen metrics for evaluating performance of the different experiments were R^2 and the root mean squared error (RMSE). In addition, K -Fold cross validation was performed with $K = 5$ to better estimate the out of sample performance. The splits were stratified across the different neighborhoods and times of the day to ensure a proportional distribution of the observations across space and time among the splits.

The reported error metrics in the following sections consist of the mean of such metrics across all folds and individual neighborhood models. For Experiments 6 and 7 (see Table 3), metrics were also averaged across the five random subsets created for each experiment. The exact equation matches precisely Equation 3a in Schmidt's research (7).

Experiments

The chosen data, model, and evaluation approach have been detailed in the previous sections. These different components were combined in several ways to obtain data to answer the proposed questions. Table 3 describes the different experiments that were performed.

Results

The R^2 and RMSE metrics for Experiments 1 to 8 are presented in Figure 2. All experiments performed better

Table 2. Traffic Count Data Description

Columns	Description
Camera ID	Unique ID for the mobility camera
Date	Corresponding date for the data point
Hour (h)	Corresponding hour for the data point
Traffic count $\in [h - 1, h)$	Number of cars observed by the camera between the indicated hour and the previous one

Table 3. Summary of Experiments

No.	Features used	Description	Purpose
1	None	Baseline model consisting of the mean occupancy of each neighborhood over all time periods as the prediction value.	Basic baseline for comparison with most advanced models.
2	Time and NPR	Reproduction of the experiments by Schmidt (7) with same data, preprocessing, and model.	Obtain metrics for previously used features in literature to compare with traffic data results.
3	Traffic—1 h	Models only use 1 h of traffic data as features.	Assess whether traffic data are useful predictors.
4	Traffic—12 surrounding hours	Models use 6 h before and after the target time step as features.	Assess whether patterns over longer periods of time hold more information.
5	Traffic—12 previous hours	Models use 12 h before the target time step as features.	Assess potential for forecasting.
6	Traffic—50 random cameras groups	Models use 6 h before and after the target time step as features, but only for 50 cameras across the city.	Study the effects of the quantity of cameras on performance.
7	Traffic—10 random cameras groups	Models use 6 h before and after the target time step as features, but only for 10 cameras across the city.	Study the effects of the quantity of cameras on performance.
8	Time, NPR and traffic—12 surrounding hours	Models use both the features proposed by Schmidt (7) as well as the most complex view of traffic data.	Determine whether patterns between features proposed by Schmidt (7) and traffic hold more predictive power together than individually.
9	Traffic—random groups in the north, west, south, and east	Models use 6 h before and after the target time step as features, but only for selections of cameras that are close in space and away from the center.	Study the effects of the spatial dimension on performance.

Note: NPR = National Parking Database.

than the baseline when looking at median R^2 and RMSE. The results obtained in Experiment 2, which replicated the methodology by Schmidt (7), obtained slightly worse results than the ones reported in that work. A potential explanation for that is the difference in time spans for both works, and thus the size of the data available for training. Whereas this paper worked with 8 months of data, Schmidt used three full years. This means that the model trained by Schmidt had the opportunity to spot seasonality patterns, and also simply had more data for it to learn the relationship between the features and target.

Traffic features improved the performance with respect to the baseline in all cases. The view of traffic holding 12 surrounding hours of data (Experiment 4) outperformed the one using only 1 h (Experiment 3), as well as Schmidt's (7) approach (Experiment 2). Using 12 previous hours of traffic data before the predicted time step (Experiment 5)

instead of surrounding it (Experiment 4) delivered very close although slightly worse results.

Using random groups of 50 cameras and 12 surrounding hours of traffic data (Experiment 6) performed slightly worse than using all 164 cameras (Experiment 4). The experiments with 10 cameras (Experiment 7) did even worse, but were still clearly superior to using all cameras, but only 1 h of data (Experiment 3).

Using both the features proposed by Schmidt (7) and 12 surrounding hours of traffic data simultaneously (Experiment 8) resulted in the best performance across all experiments, clearly outperforming using the different features in isolation.

The results for Experiment 9 can be found in Figure 3. For each group of cameras, the performance metrics are grouped by *stadsdeel* (or districts), an administrative division of Amsterdam above neighborhoods in the

Figure 2. Distribution of R^2 and root mean squared error (RMSE) per neighborhood for Experiments 1 to 8.

Figure 3. R^2 scores for models trained with 10 cameras, either randomly selected or geographically selected, across different areas.

hierarchy. No clear relationship was found between the location where traffic data were obtained and the error differences across areas in the city.

Figure 4. Mean R^2 across folds for every neighborhood as obtained in Experiment 8.

Looking closer to the individual results per neighborhood revealed that performance obtained across neighborhoods had significant differences. Figure 4 shows a map with the mean R^2 across folds for every neighborhood as obtained in Experiment 8. This behavior was also observed by Schmidt (7). Central neighborhoods seem to perform worse than those further away, but neighborhoods with high and low R^2 scores can be found

in both areas. No clear patterns in data that could potentially explain this behavior were found while exploring the results of the experiments, and whether this could be because of trips through the center or a special behavior in space demand remains unclear.

The examination of typical parking occupancy in different neighborhoods also revealed that most neighborhoods are well below the parking pressure recommended by Shoup (1). Central neighborhoods showed a mean occupancy of 68% and a standard deviation of 16%, whereas other neighborhoods showed a lower mean occupancy of 53% and a standard deviation of 13%.

Discussion

In this section, a more in-depth analysis of the obtained results can be found, with special attention to the research questions presented in the introduction.

Traffic Counts as a Predictive Feature

The first research question posed by this paper was, Does sensor-gathered data about traffic flows reduce prediction error on existing imputation and forecasting models for on-street parking occupancy in the city of Amsterdam? The results from the executed experiments, where in all cases the baseline model was outperformed in both R^2 and RMSE, showed that the generated traffic features were indeed useful in both imputing and forecasting parking occupancy in the case of the city of Amsterdam. Furthermore, the results of Experiment 8 outperforming the results of Experiment 2 specifically confirmed that using traffic features on top of the features already proposed by Schmidt (7) reduced the obtained error. Results also showed that representations of the state of traffic that describe longer time spans performed better than simpler ones, as seen in Experiments 3 and 4.

This paper has examined both imputation and forecasting for parking occupancy. Experiment 5 proposed only using traffic data before the time step that was to be predicted, effectively turning the problem into one of forecasting, and performance was on a par with the imputation scenarios. With these results in hand, we can state that traffic counts could be used in both imputation and forecasting scenarios with positive outcomes.

Number of Traffic Counting Points and Performance

The next proposed research question was, How does the number of sensors affect the prediction error? Experiments 6 and 7 limited the number of mobility cameras used from the original 164 to 50 and 10, respectively. The decrease in the number of used cameras

Figure 5. Distribution of correlation coefficients between traffic counts of 164 cameras.

gradually decreased predictive performance in relation to R^2 and RMSE, but the results were still clearly superior to the baseline and also to the features proposed by Schmidt (7). This indicated that a small number of vehicle counting-points could be enough for the purpose of predicting parking occupancy, and that increasing the number of counting points may not linearly improve performance.

Two hypotheses could explain this behavior. One is that the high correlation between the traffic counts obtained by the different cameras translated into very little information being lost when cameras were discarded. Figure 5 shows the distribution of the correlation coefficients between all cameras, which clearly indicates that most cameras have a strong positive correlation with the others. The other hypothesis would be that the autoencoding procedure used in the methodology of this paper might have produced a summarizing effect on the traffic counts, which would have weakened the signals from the more complex views of traffic.

Effects of Camera Positioning on Performance

The last proposed research question is, How does the geographical position of sensors affect the prediction error? The results from Experiment 9 did not show any clear relationship between the positioning of cameras and model performance across the spatial dimension. The intuitive idea that a camera would help more in predicting areas close to it and less in predicting areas far away in the city cannot be validated with the obtained results.

There are several ideas that could explain this behavior. First of all, Amsterdam is far from being a large city in relation to area. With optimal traffic conditions, a vehicle could drive from the west to the east and back in

less than an hour. This might translate in the area of influence of the signal given by the cameras to be large enough to cover most of the city area, thus effectively making any camera a useful source of data for predicting across the entire city. To put this to the test, a similar methodology could be reproduced in larger cities. The second possible explanation, as in the previous discussion on the impact of the quantity of sensors, is that the traffic counts from the cameras across Amsterdam are highly correlated (see Figure 5). This is not surprising given that traffic generally exhibits a strongly periodic temporal pattern with the time of the day.

Although it has not been explored in this work, we believe that further steps could be taken to enable the predictive model to learn from geographical features of the data. A possible way to do this would be to include data on the position of the counting cameras and the neighborhoods. We believe that by doing this, there is a possibility that the model would be able to learn complex geographical patterns in the parking and traffic spatial-temporal data that could improve its predictive performance. Further work would be required to validate this hypothesis.

Conclusions

This paper has experimented with traffic flow data in imputation and forecasting scenarios for on-street parking occupancy prediction, with the intent of assessing whether such data are useful predictive features. With the use of several data sources containing real data for the city of Amsterdam, empirical results indicated that traffic flow data are valid for imputing and forecasting parking occupancy. The performance obtained in Amsterdam through the use of traffic data outperformed both the proposed baseline and the previously proposed features found in literature. Furthermore, a combination of traffic data with such features led to the best results across all experiments. The quantity and positioning of cameras were examined and results indicated that a relatively small number (i.e., 10) of cameras delivered performance results close to using all 164 available cameras, and that the geographical placement did not lead to relevant differences in performance. This could be good news for cities where sensor infrastructure is scarce or unbalanced across space. Similar results should be expected of any other vehicle-counting sensor, such as induction loops or radar systems.

Although care should be taken when translating results to other cities whose spatial characteristics may differ significantly from each other, the results of this work do concur with those obtained by Ziat et al. (21) and Yang et al. (22) in Lyon and Pittsburgh, thereby

contributing more evidence that traffic data could be used to successfully estimate parking occupancy.

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: P. Martín Calvo, B. Schotten, E. Dugundji; data collection: P. Martín Calvo, B. Schotten; analysis and interpretation of results: P. Martín Calvo, B. Schotten, E. Dugundji; draft manuscript preparation: P. Martín Calvo. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Pablo Martín Calvo <https://orcid.org/0000-0001-6170-2573>

Bas Schotten <https://orcid.org/0000-0001-9626-2088>

Elenna R. Dugundji <https://orcid.org/0000-0001-9277-9485>

References

1. Shoup, D. C. Cruising for Parking. *Transport Policy*, Vol. 13, No. 6, pp. 479–486. <https://doi.org/10.1016/j.tranpol.2006.05.005>.
2. Van Ommeren, J., D. Wentink, and J. Dekkers. The Real Price of Parking Policy. *Journal of Urban Economics*, Vol. 70, No. 1, 2011, pp. 25–31. <https://doi.org/10.1016/j.jue.2011.02.001>.
3. Mathur, S., T. Jin, N. Kasturirangan, J. Chandrashekhara, W. Xue, M. Gruteser, and W. Trappe. ParkNet: Drive-By Sensing of Road-Side Parking Statistics. *Proc., MobiSys'10 8th International Conference on Mobile Systems, Applications, and Services*, San Francisco, CA, 2010.
4. Bock, F., S. Di Martino, and M. Sester. What are the Potentialities of Crowdsourcing for Dynamic Maps of On-Street Parking Spaces? *Proc., 9th ACM SIGSPATIAL International Workshop on Computational Transportation Science*, Burlingame, CA, 2016. DOI: 10.1145/3003965.3003973.
5. Bock, F., S. Di Martino, and A. Origlia. Smart Parking: Using a Crowd of Taxis to Sense On-Street Parking Space Availability. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 21, No. 2, pp. 496–508. <https://doi.org/10.1109/TITS.2019.2899149>.
6. Pierce, G., and D. Shoup. Getting the Prices Right: An Evaluation of Pricing Parking by Demand in San Fran. *Journal of the American Planning Association*, Vol. 79, No. 1, 2013, pp. 67–81.

7. Schmidt, J. *Imputing Parking Usage on Sparsely Monitored Areas Within Amsterdam Through the Application of Machine Learning*. Master's thesis, Universiteit van Amsterdam, 2019. <https://scripties.uba.uva.nl/search?id=690472>.
8. Arnott, R., and E. Inci. An Integrated Model of Downtown Parking and Traffic Congestion. *Journal of Urban Economics*, Vol. 60, No. 3, 2006, pp. 418–442. <https://doi.org/10.1016/j.jue.2006.04.004>.
9. Calthrop, E., and S. Proost. Regulating On-Street Parking. *Regional Science and Urban Economics*, Vol. 36, No. 1, 2006, pp. 29–48. <https://doi.org/10.1016/j.regsciurbeco.2005.04.002>.
10. Fosgerau, M., and A. de Palma. The Dynamics of Urban Traffic Congestion and the Price of Parking. *Journal of Public Economics*, Vol. 105, 2013, pp. 106–115. <https://doi.org/10.1016/j.jpubeco.2013.06.008>.
11. Vickrey, W. The Economizing of Curb Parking Space. *Traffic Engineering*, Vol. 25, No. 2, 1954, pp. 62–67.
12. Van Ommeren, J. N., D. Wentink, and P. Rietveld. Empirical Evidence on Cruising for Parking. *Transportation Research Part A: Policy and Practice*, Vol. 46, No. 1, 2012, pp. 123–130. <https://doi.org/10.1016/j.tra.2011.09.011>.
13. Bock, F., and M. Sester. Improving Parking Availability Maps using Information from Nearby Roads. *Transportation Research Procedia*, Vol. 19, 2016, pp. 201–214. <https://doi.org/10.1016/j.trpro.2016.12.081>.
14. Ionita, A., A. Pomp, M. Cochez, T. Meisen, and S. Decker. Where to Park? Predicting Free Parking Spots in Unmonitored City Areas. *Proc., 8th International Conference on Web Intelligence, Mining and Semantics*, Novi Sad Serbia, 2018.
15. Pullola, S., P. K. Atrey, and A. E. Saddik. Towards an Intelligent GPS-Based Vehicle Navigation System for finding Street Parking Lots. *Proc., 2007 IEEE International Conference on Signal Processing and Communications*, Dubai, 2007.
16. Ji, Y., D. Tang, P. Blythe, W. Guo, and W. Wang. Short-Term Forecasting of Available Parking Space Using Wavelet Neural Network Model. *IET Intelligent Transport Systems*, Vol. 9, No. 2, 2015, pp. 202–209. <https://doi.org/10.1049/iet-its.2013.0184>.
17. Zheng, Y., S. Rajasegarar, and C. Leckie. Parking Availability Prediction for Sensor-Enabled Car Parks in Smart Cities. *Proc., IEEE 10th International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, Singapore, 2015.
18. Rajabioun, T., and P. Ioannou. On-Street and Off-Street Parking Availability Prediction Using Multivariate Spatio-temporal Models. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 16, No. 5, 2015, pp. 2913–2924. <https://doi.org/10.1109/TITS.2015.2428705>.
19. Badii, C., P. Nesi, and I. Paoli. Predicting Available Parking Slots on Critical and Regular Services by Exploiting a Range of Open Data. *IEEE Access*, Vol. 6, pp. 2018, pp. 44059–44071. <https://doi.org/10.1109/ACCESS.2018.2864157>.
20. Fan, J., Q. Hu, Y. Xu, and Z. Tang. Predicting the Vacant Parking Space Availability: An LSTM Approach. *IEEE Intelligent Transportation Systems Magazine*. 2020. <https://doi.org/10.1109/MITS.2020.3014131>.
21. Ziat, A., B. Leroy, N. Baskiotis, and L. Denoyer. Joint prediction of Road-Traffic and Parking Occupancy over a City with Representation Learning. *Proc., IEEE Conference on Intelligent Transportation Systems*, Rio de Janeiro, Brazil, 2016.
22. Yang, S., W. Ma, X. Pi, and S. Qian. A Deep Learning Approach to Real-Time Parking Occupancy Prediction in Transportation Networks Incorporating Multiple Spatio-Temporal Data Sources. *Transportation Research Part C: Emerging Technologies*, Vol. 107, 2019, pp. 248–265. <https://doi.org/10.1016/j.trc.2019.08.010>.
23. Hössinger, R., K. Heimbuchner, and T. Uhlmann. Development of a Real-Time Model of the Utilisation of Short-Term Parking Zones. *Proc., 19th Intelligent Transport Systems World Congress*, Vienna, Austria, 2012.
24. Gemeente Amsterdam. Amsterdam Neighbourhood Map, 2020. https://data.amsterdam.nl/data/?modus=kaart&achtergrond=topo_rd_zw¢er=52.3704876%2C4.8934422&lagen=gebind-buurt%3A1&zoom=7.
25. Bellman, R. R. E. *Dynamic Programming*. Princeton University Press, NJ, 1957.
26. Friedman, J. H. Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, 2001, pp. 1189–1232. <https://doi.org/10.1214/aos/1013203451>.
27. Schapire, R. E. A Brief Introduction to Boosting. *Proc., 16th International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, 1999.
28. Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay. Scikit-Learn: Machine Learning in Python. *Journal of Machine Learning Research*, Vol. 12, 2011, pp. 2825–2830.