

Where to Park? Predicting Free Parking Spots in Unmonitored City Areas

Andrei Ionita

Informatik, RWTH Aachen University
Aachen, Germany
andrei.ionita@rwth-aachen.de

André Pomp

Institute of Information Management
in Mechanical Engineering, RWTH
Aachen
Aachen, Germany
andre.pomp@ima.rwth-aachen.de

Michael Cochez

Fraunhofer FIT
Aachen, Germany
Computer Science 5, RWTH Aachen
Aachen, Germany
michael.cochez@fit.fraunhofer.de

Tobias Meisen

Institute of Information Management
in Mechanical Engineering, RWTH
Aachen
Aachen, Germany
tobias.meisen@ima.rwth-aachen.de

Stefan Decker

Computer Science 5, RWTH Aachen
Aachen, Germany
Fraunhofer FIT
Aachen, Germany
stefan.decker@dbis.rwth-aachen.de

ABSTRACT

Several smart cities around the world have begun monitoring parking areas in order to estimate free spots and help drivers that are looking for parking. The current results are indeed promising, however, this approach is limited by the high costs of sensors that need to be installed throughout the city in order to achieve an accurate estimation rate. This work investigates the extension of estimating parking information from areas equipped with sensors to areas that are missing them. To this end, similarity values between city neighborhoods are computed based on background data, i.e., from geographic information systems. Using the derived similarity values, we analyze the adaptation of occupancy rates from monitored to unmonitored parking areas.

CCS CONCEPTS

• **Information systems** → *Clustering; Semantic web description languages*; • **Computing methodologies** → *Cluster analysis*;

KEYWORDS

smart parking, machine learning, semantic annotation, data mining

ACM Reference Format:

Andrei Ionita, André Pomp, Michael Cochez, Tobias Meisen, and Stefan Decker. 2018. Where to Park? Predicting Free Parking Spots in Unmonitored City Areas. In *WIMS '18: 8th International Conference on Web Intelligence, Mining and Semantics, June 25–27, 2018, Novi Sad, Serbia*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3227609.3227648>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WIMS '18, June 25–27, 2018, Novi Sad, Serbia

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5489-9/18/06...\$15.00
<https://doi.org/10.1145/3227609.3227648>

1 INTRODUCTION

Parking is a known problem in cities. Worldwide we are experiencing an increase in the number of cars [21]. Currently, about 30% of the traffic in cities is caused by cars that are actively searching for parking [20]. Drivers often end up double-parking their cars, which blocks other cars, thus causing unneeded stress. Drivers that are circling for a parking space may cause safety issues, as they are often distracted and are not paying attention to cyclists and pedestrians. When circling for parking spaces, additional fuel is consumed, which has an economical as well as an environmental impact.

Existing parking spaces can be more efficiently occupied if the drivers know about their availability. This information needs to be available in advance, so that drivers can take the decision to drive towards a highly probable free parking space early enough and not get stuck in a traffic bottleneck. A software system predicting available spaces would, ideally, take into account factors such as current parking space availability, traffic, events in the near vicinity, weather, and many more. Existing forecasting systems often start by collecting statistics on available parking spaces. Usually, mounted sensors observe when a car occupies and leaves a parking space. Acquiring the required data, however, be it on the parking spaces themselves or the complementary information, is, most of the time, the bottleneck for these approaches.

Our approach introduces the notion of a *parking demand profile*, which reflects the time of the day where parking occurs and its duration for a given area. Furthermore, we examine how machine learning models for parking occupancy can be transferred between city areas with a similar profile. We start by giving an intuitive example which bridges the understanding path towards the rest of the work. Further, other contributions to the field of smart parking are presented, followed by the concrete approach. The evaluation in this case carries out the instantiation of the proposed approach with an actual use case. We present the results and an analysis. We finish off by pointing out aspects that can be further pursued and by summarizing the outcome. This work is a short version of an earlier master thesis by the first author of this article [10].

2 MOTIVATING EXAMPLE

In this section, we provide a motivating example illustrating the necessity for introducing the concept of *parking demand profiles* to improve the accuracy of machine learning models for predicting free parking spots in smart cities.

The scenario consists of a large smart city with different parking areas, which are either dedicated parking spots (e.g., car parks or parking lots) or parking spaces on the side of the street. Some of these parking areas are equipped with sensors to measure the occupancy rate whereas other areas are tracked using cameras or parking meters. Due to the high costs for tracking all parking areas (especially those which offer free parking), most parking areas are not tracked at all. To improve the parking situation in the city, the local government decides to publish the available data sources, such as the sensor measurements, on their Open Data platform.

Data scientists can now obtain the data and start building predictive models for parking areas at which data are available. For that, the data scientist examines the data, performs feature engineering and selection, trains the models and develops an application that predicts the occupancy rates for different locations of the city.

Examining this scenario, we identify several drawbacks. Due to the heterogeneity of available parking data sources and information, it becomes quite hard for the data scientist to build the models. First, the data scientist has to understand each data source in detail based on the information provided on the Open Data platform. If the data were provided homogeneously (e.g., by using RDF and a well-defined ontology), it would become easier for the data scientist to understand the data from one source and to aggregate the data from multiple sources. In addition, the application can be used in other cities, which provide their data based on the same ontology.

Besides the understanding of the data, the data scientist still faces the problem of not covering all parking areas of the city. If we assume that different models for the different parking areas have been built, it is unclear whether any of these models can be applied on untracked areas. Hence, our approach exploits the use of parking profiles which represent parking in a city area. One example would be an area consisting of office buildings. Here, parking demand is usually high during working time, e.g., between 8 and 18 o'clock. Two such areas, perhaps in different parts of a city, or even in different cities, likely have a similar parking situation, i.e., between 8 and 18 on weekdays there is a high parking demand. In case of restaurants, on the other hand, we see a spike in parking demand in the evening, usually from 18 until 22, and even more so on weekends. In residential areas, the cars are parked in the evening and leave again early in the morning. A measure that would capture parking demand will therefore be based on the stay duration of customers or employees of the particular services, the available number of shops and restaurants or other point of interests (e.g., event halls).

For such city areas (e.g., office buildings, restaurants, residential areas, etc.) that have parking occupancy data, the data scientist can build the predictive models. Then, using parking profiles, one can attempt using models created for tracked areas in untracked ones. Hence, we investigate in this article how we can transfer these models to other city areas for which no parking occupancy data is available, but that have similar parking demand profiles.

3 RELATED WORK

The goal of improving the parking situation by estimating free parking spots is well known for many decades. Hence, there are many works related to the problem we are addressing.

Xu et al. [23] make real-time parking availability estimations based on data collected from mobile phones. The system uses algorithms incorporating statistical weighted schemes and Kalman filters. Additionally, the authors create parking availability profiles based on historical data and statistical algorithms.

Other works making use of smart phones are the ones by Nandugudi et al. [14], Koster et al. [11], and Chen et al. [7]. After data collection, often using the sensors in the smartphone, they use different machine learning models, such as probabilistic distribution, Bayesian, and Hidden Markov models, as well as fuzzy logic to predict the available parking spaces.

Instead of using mobile devices that are carried by drivers, ParkNet, developed by Mathur et al. [13] is a system made up of dedicated vehicles that captures parking space information while driving. Every ParkNet vehicle is equipped with a GPS receiver and an ultrasonic sensor facing sideways. The latter determines whether it passes by parking spaces and whether they are occupied. The data is sent to a central server that aggregates it, in order to build parking space occupancy maps in real-time. The information is queried by clients that search for a free parking space. The system was evaluated in Highland Park, New Jersey and San Francisco on 500 miles road-side parking data and yield 95% accurate parking maps and 90% parking occupancy accuracy.

Compared to approaches that collect data via mobile phones or dedicated cars, other approaches use data collected by sensors, parking meters or cameras that are installed on streets or parking lots. For instance, Tiedemann et al. [22] developed a prediction system that gives estimated occupancies for parking spaces in Berlin, Germany. The occupancy data is collected online via roadside parking sensors and the prediction is realized using neural gas machine learning combined with data threads. The authors notice that some factors play a significant part in the predictions, such as holidays, weather and thus, they use the neural gas clustering to separate the data, before the data thread method is applied. This perfectly matches our idea of considering further information about the circumstances in a city. However, our approach considers more detailed information about specific districts rather than general information like holidays, which usually hold for an entire city.

Similar to the approach presented by Tiedemann et al. [22], Hössinger et al. [9] present a simple real-time occupancy model based on various pieces of data collected in the city of Vienna, Austria. An average day curve model was built using the ticket data from mobile phone parking, the counts of car parks, and the traffic flow volumes in the city. The data was collected following agreements with the respective mobile phone companies, through surveys and by accessing a dedicated traffic website, respectively. The predictions are valid for short-time spans and applicable to the above mentioned city.

Located in Spain (Barcelona), Caicedo et al. [4] developed a methodology for predicting real-time parking space availability. The probabilistic algorithm consists of three subroutines: allocating simulated parking requests, estimating future departures, and

forecasting parking availability. The forecast has been reported to improve as the system registers arrivals and departures. Similar to our work, Caicedo et al. [4] take further factors into account, such as duration of stay and capacity of every operating parking facility.

Rajabioun and Ioannou [17] introduce an information system for parking guidance that enables communication between vehicles and the infrastructure. It proposes a prediction algorithm that forecasts the availability for parking locations based on real-time parking information. It takes into account parameters such as parking duration, arrival time, destination, pricing, walking distance, parking capacity, etc. Their algorithm uses a probabilistic density distribution model. The parking data was collected both from on-street parking meters and off-street garages in Los Angeles and San Francisco, USA. In a follow-up paper, Rajabioun and Ioannou [18] propose a multivariate autoregressive model that considers the temporal and spatial correlations of parking availability when making predictions. The authors hold that the model, which is integrated into a parking guidance and information system, recommends parking locations with high accuracy.

In contrast to Rajabioun and Ioannou [17], Chen [6] tackles the parking problem by aggregating parking lots. Their findings show that the prediction error of parking occupancy decreases by combining multiple parking lots. The trained models take into account factors such as day, time, event, distance, parking price, etc. The author evaluated multiple models, such as ARIMA, linear regression, support vector regression and feed forward neural network. It turns out that the neural networks algorithm scores the best when the model is evaluated on the *SFpark* data [2]. The findings of Chen confirm our idea building models for smaller areas, such as one district and one parking lot. However, our work goes one step further by additionally transferring the trained model to districts where no data is available.

Richter et al. [19] address the parking prediction problem with the focus on model storage in vehicles. The authors train models of various granularity that would predict parking availability based on the information contained: a one-day model per road segment, a three-day model per road segment, and a seven-day model per road segment. Additionally, models based on regions and time intervals computed by clustering are tried out. Hierarchical clustering with complete linkage is employed. The models are evaluated on street data from the *SFpark* project [2] and reach a prediction success rate of about 70%.

While all the presented approaches tackle the problem of improving the parking situation in cities with different ideas and implementations, they fundamentally differ from our work. Some approaches match with our idea of adding additional information about the current circumstances in a city or by building models for smaller areas. However, none of the approaches focuses on building a solution that can be ported to other cities. As opposed to this, our presented approach uses a public pre-defined ontology and data in RDF format as input in order to calculate the prediction models and similarity measures. Further, none of the presented approaches aims at training prediction models for small areas where data are available and transfer those models to other areas where no data is available. Our approach aims exactly at this goal for enabling a more accurate prediction in districts or streets where no parking data are available.

4 SELECTED DATA SOURCE

To implement and evaluate our idea, we use the data from the *SFpark* project. This project was realized by the San Francisco Municipal Transportation Agency (SFMTA), the city agency that manages the city's transportation, which includes on-street parking [1] [3]. The SFMTA establishes parking rates for on-street parking meters. Before the project started, parking rates were the same all day, every day, independent of the parking demand. By implementing a demand responsive pricing scheme, parking availability improved dramatically.

In conducting the project, nine pilot areas were chosen for monitoring. Out of these areas, seven were selected to have new pricing policies, while two were control areas. The number of metered spaces used was 6000, which amounts for 25% of the city's total. The meters allow rates to be deployed remotely and they transmit data to a central server through a wireless connection.

The data was collected using parking sensors. These provided the central server with the information needed to calculate the demand-responsive parking rates and provided real-time parking availability information. A parking sensor is a magnetometer that detects changes in the earth's electromagnetic field. A total of 11700 sensors were deployed, resulting in 8000 spaces that were equipped with one or two sensors. The sensors delivered valid data from April, 2011 to December, 2013. The sensors can suffer from environmental noise, such as electromagnetic interference, early battery degradation or street construction.

SFpark made available real-time information on parking rates and parking occupancy through a smart phone application. The *SFpark* project and its success played an important role when choosing to base our project on it.

5 CONCEPT

For implementing our approach, we extract parking data from *SFpark*. In addition, we collect city data from various sources (cf. Section 5.2). To ensure that our approach will be compatible with data available in other cities, we annotate our data using ontologies from the CityPulse project [16] and store them in RDF format. Afterwards, we merge the data sets and perform a clustering process. Next, we define the parking profile by introducing two similarity functions: *cosine similarity* and *earth mover's distance*. Finally, we explain the machine learning training process and define the estimations for clusters without parking data.

5.1 Parking Data

We consider the following types of data as parking data: *parking occupancy* contains information on the availability of parking spaces; *traffic data* contains information regarding the city traffic, which is relevant for parking; *weather data* contains weather information for the same area as for the parking problem; *event data* contains event information which may have an impact on parking; *parking revenue data* contains economic information on parking, whose relevance may influence parking prediction. *fuel price data* contains prices of fuel in the region for which we build the models. Each piece of data is geographically referenced by a *location unit*, e.g., street block, district or city. An overview of the different properties available in the data set are shown in Table 1.

Parking Occupancy		Traffic	
date and time	Recorded usually at full hours or in periodic time intervals	date and time	recorded usually at full hours or in periodic time intervals
parking capacity	The total number of parking spaces at the given location	traffic value	typically expressed as average traffic road occupancy, average vehicle count, median speed, or average speed of the traveling cars
parking price	The price of a ticket at the certain location and the given time in a given currency		
parking occupancy	Expressed either as rate (subunitary fraction or percent) or in absolute numbers		
Events		Weather	
event name class	the name of the event and its class (road closure, rise of parking demand)	temperature	may be current temperatures or maximum and minimum values per day
		precipitation	expressing the quantity of rain or snow for the corresponding time interval
Fuel Price		Parking revenue	
type of fuel	gasoline, diesel, etc.	payment type	the way the driver opted to pay for parking, e.g., cash, credit card, etc.
price per unit	provided as the price per liter or per gallon.	paid amount	the amount in US dollars, Euro or other currency

Table 1: An overview of the properties available in the data used from the SFpark project. Each of these data types also has the location unit id, as well as the date and time (interval) when it occurred or was measured. In some cases, the location is a somewhat larger or smaller area as the unit. The time information is provided in different granularities (e.g., per minute, per hour, per day, etc.).

5.2 City Data

We established that city data reflects parking demand in a city area. We obtain it by collecting public amenity information as provided by OpenStreetMap¹. The OSM data is available as shapefiles containing geometries such as points, polylines and polygons. We extract the *points of interest* (POIs). A POI contains an *amenity* attribute indicating the public service located at this position as it was annotated by the OSM users (cf. Figure 1).

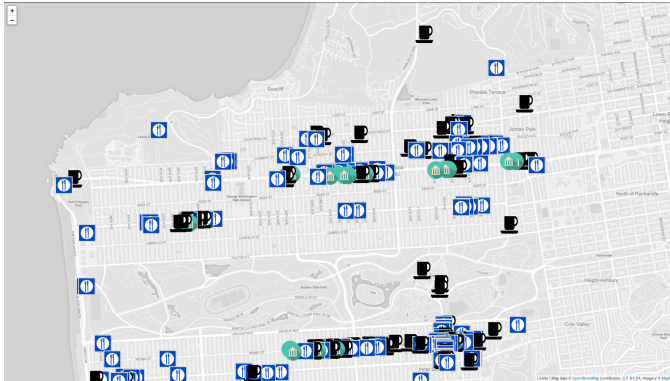


Figure 1: A map indicating public amenities (cafes, restaurants, banks) found at points of interest in OSM.

¹<https://www.openstreetmap.org> The maps used in this article are ©OpenStreetMap contributors

The amenity information accounts for the times of parking (i.e., morning, day, night, etc.) and, up to some degree, for the duration of parking. An example of a service indicating this is Google Maps. It displays typical *visiting duration* values and popularity of the place for specific points in time. The average values are based on the users' smart phone GPS sensors (cf. Figure 2). To obtain the duration values, we manually extract information from Google Maps. The duration information is aggregated by Google using a crowdsourcing approach.




Figure 2: An example of visiting duration information found on Google Maps.

5.3 RDF Annotation

We emphasized that our approach should be relevant for any city that wants to solve its parking problem. Thus, all input data used in the present approach is in RDF format, in order to establish a common format that can be used for other cases. Since the actual parking and city data is only available as raw values, we need to annotate it as RDF in the first place. The default process involves the extraction of data, which is afterwards available for further processing. To annotate the city and parking data, we use Apache Jena. As underlying ontologies, we reuse those created as part of the CityPulse project [16]. Afterwards, we can easily extract the data in a well-defined format using SPARQL queries.

5.4 Merging City and Park Data

In order to combine the parking and city data, both sets of data require a common location unit. The parking location units are provided together alongside the various types of parking data. The city data, on the other hand, references POI geometries, which are points expressed in a particular reference system, which differs from the one of the parking locations. Therefore, after establishing the coordinate systems of both geometries, we define a *merge distance* that matches a parking space to a public amenity.

The merge distance can be intuitively understood as the radius around a public amenity. It is defined to represent the parking area that is relevant for a particular public amenity, or, more straightforward, the walking distance from the parked car to, e.g., the restaurant, the office, the bank, etc. Concrete instances of the merging distance can be found in Section 6.3.

5.5 Clustering

In a realistic scenario, the available parking data does not cover the whole city surface. In fact, it is a fraction of it. Therefore, we first separate the area *with* parking data from the area *without* parking data. Based on this initial split, we perform clustering to further separate these regions into smaller areas. By splitting into city areas, we are making sure that smaller regions lead to more representative parking profiles and therefore parking estimations.

As we want an exclusively location-based separation, we may employ K-Means, DBSCAN or OPTICS to cluster the city areas. The distance is calculated between (latitude, longitude)-pairs of location unit coordinates corresponding to one parking unit. There are two clustering processes executed, one for the city area *with* parking data, another one for the city area *without* parking data. The number of clusters chosen in each area is kept proportional to the number of total location units that each city area contains. Since having control over the number of clusters is the goal here, we choose to use K-Means, where we provide the number of expected clusters as input. More details on the concrete value of k and an overview about the clusters can be found in Section 6.4.

5.6 Similarity Functions

To calculate the similarity of city areas, we use *cosine similarity* and *earth mover's distance*. For this, we use two representations of city areas: *cluster vectors* and *cluster Gaussians*. As data, we use the city data described in Section 5.2.

5.6.1 Cosine Similarity. To form the cluster vectors, we first divide all amenities into categories $Cat_1, Cat_2, \dots, Cat_n$. The criteria for division will be their average visiting duration. For example, a short duration category of up to 30 minutes, a medium duration between 31 and 90 minutes and a large duration of above 90 minutes stay. Each cluster gets represented by an n -dimensional vector, whose components correspond to the amenity categories. The magnitude of component i is equal to the number of amenities of category Cat_i that can be found in that particular cluster. Compare Figure 3 for a general representation.




Figure 3: An example of a cluster vector for three categories.

The cosine similarity between two vectors is defined as the cosine of the angle between the two vectors:

$$\cos(\theta) = \frac{A \cdot B}{\|A\|_2 \|B\|_2} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}} \quad (1)$$

where A_i and B_i are the components of vector A, respective B.

Unlike the earth mover's distance, the cosine similarity implementation uses the direct mathematical formula by plugging in the magnitudes of the respective vector components.

5.6.2 Earth Mover's Distance. A cluster Gaussian is a *kernel density estimation* among amenity probability distributions. In turn, an amenity probability distribution is represented as *Gaussian kernel*. To construct a cluster Gaussian, we first collect the average visiting duration and standard deviation for the individual amenities. A cluster that contains one amenity A is represented as a Gaussian curve, i.e., normal probability distribution. The curve's center is at the average duration of the amenity A and its standard deviation is the one of the amenity. When n amenities A exist in the cluster, the representation will be an A curve multiplied n times. Multiple amenities, each appearing multiple times, will result in a curve that is the linear combination of the individual representations of the amenities as normal distribution curves. Compare Figure 4 for a visualization of the summing process.

$$EMD(C_i) = \sum_{j=1}^{|amenities|} K_{ij} \times A_j \quad (2)$$

$\forall i \in \{1, \dots, |clusters|\}$ and $\forall j \in \{1, \dots, |amenities|\}$
 where A_j is an amenity that appears K_{ij} times in the cluster C_i .




Figure 4: The summing of Gaussians resulting in a cluster Gaussian.

The earth mover's distance (EMD) is a measure used in statistics that roughly expresses the difference between position and magnitude of two curves. It is best explained by regarding the curves as the hull of earth piles. For two separate earth piles, EMD computes the minimum effort of rearranging a pile so that the shape of the other pile is obtained. Moving P particles over a distance D is equal to the effort $P \times D$. A prerequisite for this operation is that the two piles need to contain the same quantity of earth.

More rigorously, the earth mover's distance is better known in mathematics under the name Wasserstein Metric. Given two normal distributions $\mu_1 = \mathcal{N}(m_1, C_1)$ and $\mu_2 = \mathcal{N}(m_2, C_2)$, where m_1 and $m_2 \in \mathbb{R}^n$ are their respective expected values and C_1 and $C_2 \in \mathbb{R}^{n \times n}$. Then, the 2-Wasserstein distance between μ_1 and μ_2 is:

$$W_2(\mu_1, \mu_2)^2 = \|m_1 - m_2\|_2^2 + \text{trace}(C_1 + C_2 - 2(C_2^{1/2} C_1 C_2^{1/2})^{1/2}) \quad (3)$$

In practice, we will not apply the Wasserstein metric directly, but rather resort to some levels of discretization. First off, a number of so-called bins is determined. Each bin represents a unit on the X axis, the same on which the visiting duration is expressed. We will take a number of buckets equal to the maximum amenity mean and add $3 \times$ the largest standard deviation to it, as it is known that within $3 \times$ standard deviation on both sides of the mean over 99% of the Gaussian sum is covered. Moreover, an offset on the X -axis equal to $3 \times$ the maximum standard deviation is used. This way, we are sure the landscape of summed Gaussians will easily fit into the number of bins.

Notice that EMD is applicable only when the sum under both Gaussian curves is equal. Therefore, all cluster Gaussians will get normalized before EMD is computed.

5.7 Machine Learning Models

The prediction of parking occupancy is realized using machine learning. We choose to explore this methodology, following the solid results machine learning models have delivered for the various

smart parking settings investigated in Section 3. A machine learning model \mathcal{M} will be trained for every cluster *with* parking data.

The training data *features* are composed of the parking data previously enumerated. We aggregate feature values around the location unit id, as on the cluster level this is irrelevant. The occupancy rate is set as the target variable. The model training and evaluation is performed in Python via the *scikit-learn* library.

During the training phase, we evaluate different machine learning approaches ranging from simple decision trees over random forests to support vector machines and gradient boosting. In addition, we make use of grid search to determine the best model parameters for the current approach. As error metric, we use root mean square error and performed a ten-fold cross-validation. Furthermore, a model is evaluated on the other clusters with parking occupancy data.

5.8 Parking Occupancy Estimations

Once all models \mathcal{M} have been built for the clusters *with* parking data, making estimations on parking occupancy in these areas is straightforward. However, we want to apply these models, on the clusters that are missing parking data. We derive the *estimation interval* for cluster C_{wout}^j based on the model of cluster C_{with}^i as follows.

For cosine similarity:

$$E(C_{with}^i, C_{wout}^j) = [\mathcal{M}(C_{with}^i) - (1 - sim_{ij}), \mathcal{M}(C_{with}^i) + (1 - sim_{ij})] \quad (4)$$

$$\text{where } sim_{ij} = sim(C_{with}^i, C_{wout}^j) \in [0, 1]$$

For earth mover's distance:

$$E(C_{with}^i, C_{wout}^j) = [\mathcal{M}(C_{with}^i) - emd_{ij}, \mathcal{M}(C_{with}^i) + emd_{ij}] \quad (5)$$

$$\text{where } emd_{ij} = emd(C_{with}^i, C_{wout}^j) \in [0, 1]$$

$$\forall i \in \{0, \dots, |C_{with}| - 1\} \text{ and } \forall j \in \{0, \dots, |C_{wout}| - 1\}$$

X is a parking data record containing feature values. The result is an *estimation interval* that “stretches” the punctual estimation into an interval depending on the similarity value. The lower the similarity value is, the larger the length of the resulting estimation interval will be.

Notice that X should be valid for both C_{with}^i and C_{wout}^j . The averaged values used for features related to the location unit id will be transferred as they are similar to those of the target cluster C_{wout}^j .

Furthermore, we define an *estimation intersection interval*, whose purpose is to narrow down the computed estimation interval. An estimation intersection interval for the clusters C_{with}^i and C_{wout}^j is computed by intersecting the *estimation intervals* that have a better similarity among the clusters with data $C_{with}^0, \dots, C_{with}^{i-1}$ and the same cluster without data C_{wout}^j .

$$EI(C_{with}^i, C_{wout}^j) = \bigcap_{k=0}^{i-1} EI(C_{with}^k, C_{wout}^j) \quad (6)$$

where

$$\text{sim}(C_{\text{with}}^k, C_{\text{wout}}^j) < \text{sim}(C_{\text{with}}^i, C_{\text{wout}}^j), k \in \{0, \dots, i-1\} \text{ for EMD} \quad (7)$$

$$\text{sim}(C_{\text{with}}^k, C_{\text{wout}}^j) > \text{sim}(C_{\text{with}}^i, C_{\text{wout}}^j), k \in \{0, \dots, i-1\} \text{ for cosine} \quad (8)$$

$$\forall i \in \{0, \dots, |C_{\text{with}}| - 1\} \text{ and } \forall j \in \{0, \dots, |C_{\text{wout}}| - 1\}$$

6 EVALUATION

For evaluating our approach, we require parking and city data. As parking data, we use those from the *SFpark* project and determine the data sets that are relevant for our purposes. As city data, we take OpenStreetMap data from San Francisco and obtain the available public amenity information. Google Places provides the visiting duration values for the amenities. Based on these data sets, we determine relevant values for the *merge distance*. Since we cannot evaluate the performance of our approach using clusters *without* data, we evaluated the performance on clusters for which data were available. Therefore, we first calculated the similarity score between a cluster c and all other clusters. Afterwards, we calculated the estimation errors when applying the model of cluster c to all other clusters.

6.1 SFpark Parking Data

The *SFpark* data are visualized in Figure 5 using a built Leaflet application. The actual *SFpark* data has some particularities. While the *occupancy data* is provided with reference to *blocks* as location units, all the other data sets use different location units. For the traffic and events data sets, the location units are street names. For parking revenue, they are districts. In case of weather and fuel price, the location reference is valid for the whole city of San Francisco.

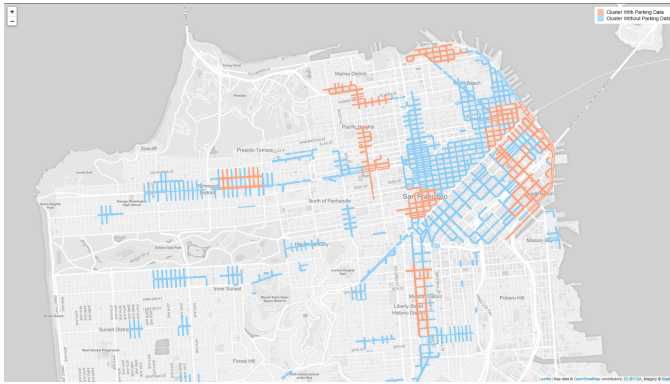


Figure 5: The blocks accounted in *SFpark*. The light blue ones are blocks *without* parking data, the light red ones are *with* parking data.

6.2 OpenStreetMap for San Francisco

Following the selection of *SFpark* data as parking data, the city data is found in the corresponding OpenStreetMap layer for San Francisco. The actual public amenity information collected from POIs is listed in Table 2.

arts_centre	dojo	marketplace	shelter
bank	embassy	music_rehearsal_place	shop
bar	fast_food	music_school	spa
biergarten	grocery	nightclub	stripclub
bureau_de_change	gym	pet_grooming_shop	studio
cafe	hookah_lounge	pharmacy	training
clinic	ice_cream	police	veterinary
clothes_store	karaoke	post_office	vintage_and_modern_resale
community_centre	lan_gaming_centre	pub	
dentist	laundry	restaurant	
doctors	library	salon	

Table 2: List of all OSM amenities found in the *SFpark* blocks.

6.3 Merging Parking and City Data

In case of *SFpark*, the blocks are given in latitude and longitude. In OpenStreetMap, the geometry is set to EPSG 4326. With both systems using the same reference, we can therefore set a *merge distance*. The distance d should express the impact that a POI P has on the block B , when $\text{dist}(P, B) = d$. For instance, it expresses the impact of the parking demand that a restaurant induces on a parking block situated at d meters away. We assign to it distances of 100m, 200m, and 400m.

6.4 Clustering

As established in Section 5.5, we apply K-Means to cluster the city areas. In the evaluation, we will refer to the number of clusters *with* parking data as *the number of clusters*. The area without parking data is going to be split into a proportional number of clusters, as the sizes of clusters should be kept roughly equal for both sides. It turns out that the proportion is approximately 2.6, following the division between the total number of blocks from each group. We have chosen two numbers of clusters to run the evaluation, namely 8 clusters and 16 clusters. The area without parking data will therefore have 20 and 41 clusters, respectively.

After running the K-Means clustering process, the Leaflet application map reveals the individual clusters by highlighting them on mouse-over. The clusters *with* parking data will turn dark red, while the clusters *without* parking data will appear in dark blue (cf. Figure 6).




Figure 6: Highlighted cluster *with* parking data on the left side and a cluster *without* parking data on the right side.

The resulting aggregated blocks is worth taking into account when training the machine learning models, as these will average over pieces of information contained in individual blocks.

6.5 Similarity Functions

Before computing *cluster vectors* and *cluster Gaussians*, we will establish the visiting duration in every amenity. For this, we use information gathered from Google Places available via Google Maps².

We manually collected information from 470 places in San Francisco, for which a maximum duration of stay was provided (the minimum duration is not always given, as indicated earlier). The data was obtained by manually navigating to every business place and writing the duration visit information in a spreadsheet. This piece of information is not accessible yet via the Google Places API³. The results are shown in Table 3 and the numbers are given in minutes and have been rounded to the nearest integer. We have included only amenities for which at least two stay duration sources were found.

Alongside this information, we need the amenity categories in order to derive the *cluster vectors*. As defined in Section 5.6.1, the categories are based on the visiting duration mean. We split them in three categories: under half an hour, 31 to 90 minutes and more than 1.5 hours. The assigned partitions for every amenity are shown in Table 3.

The calculation of *cluster Gaussians* relies on both the mean and standard deviation of the amenity visiting duration, as defined in Section 5.6.2.

6.6 Model Training and Evaluation for Clusters with Parking Data

We use four methods to train models for the clusters *with* parking data: *decision trees*, *support vector machines*, *multilayer perceptrons*, *gradient boosted trees*. As training data, the SFpark *occupancy data* is used with *street blocks* as location unit. It turns out that training on the additional SFpark data, i.e., traffic and events, encounters, some problems.

The *traffic data* do not share the same location unit with the parking occupancy's street block. However, aggregating traffic data on the district level, which is available for the occupancy data as well, does not provide an additional value to the training. For the SFpark *events data*, we encounter the same problem as for the traffic data: the location unit does not match the block. In fact, the events are marked for streets, whose association to blocks is not determinable. *Parking revenue data* is provided for districts, which again are too general to make a difference in training. Finally, *weather data* and *fuel data* are given per city, hence making even less an impact to improve the model.

As indicated in Section 5.7, the training data is averaged on all blocks so that it can be applied later on other clusters. The averaging is performed per timestamp, i.e., if multiple blocks have an occupancy record for the same time and block, the occupancy rate will be averaged for both of these. Features such as *price* and *parking capacity per block* are averaged as well. This means that the original collection of data records shrinks. In Table 4, the number

of training records per cluster is shown, together with the number of occupancy records for that cluster. On average, there are 10.5 blocks that share the same data points.

In Figure 7, a screenshot of the application showing a sample of the results is illustrated. Figure 8 displays the presented table in more detail.

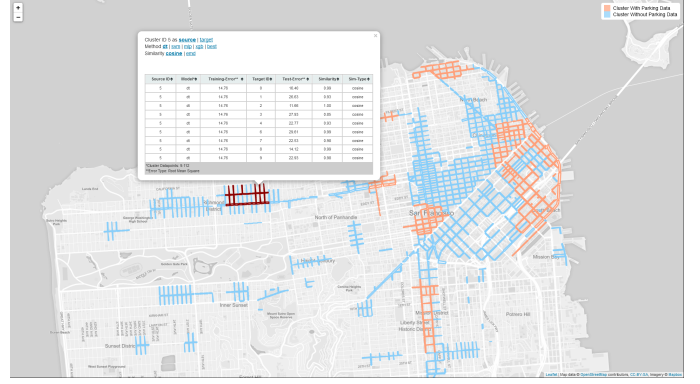


Figure 7: Selected cluster *with* parking data and the pop-up table in the Leaflet application.

6.7 Clustered- vs. Total Models

One assumption of our approach is that models originating from smaller clusters are better at predicting occupancy than models trained with the entire city area. We have compared the two types of models during our tests. For each target cluster C_{target} , we determine the source cluster C_{source} whose model has the best estimation error when applied on C_{target} . We also train a model containing the entire city area with parking data \mathcal{A} minus C_{target} 's data and compute this model's estimation error on C_{target} . As we can observe from Table 5, in the case of 8 clusters, and in Table 6 for 16 clusters, the cluster's models estimations are superior to the ones of the total model with very few exceptions.

6.8 Best Model Method

Models were trained using four methods: decision trees, support vector machines, multilayer perceptrons, gradient boosted trees. Table 7 shows the distribution of best machine learning methods in case of 8 and 16 clusters. The values were obtained by summing up the number of times a method produced the least estimation error, i.e., RMSE, among the four methods for all combinations of clusters with parking data (C_{source}, C_{target}). Extreme gradient boosting claims the first spot in both cases.

6.9 Similarity Values vs. Estimation Errors

The goal of our approach was to replace occupancy estimations for clusters where no parking data is available with estimations based on cluster similarity values. Among clusters *with* parking data, the real occupancy values are known. This enables us to compute estimation errors for cluster models, which can then be correlated with the similarity values between clusters. We use two correlation

²<http://maps.google.com>

³Google Feature Request: <https://issuetracker.google.com/issues/35827350>

amenity name	mean	stdev	cat	amenity name	mean	stdev	cat
arts_centre	110	37	3	laundry	78	16	2
bank	42	65	2	library	83	13	2
bar	121	38	3	music_school	120	30	3
cafe	76	39	2	nightclub	189	20	3
clinic	100	29	3	pharmacy	25	20	1
clothes_store	41	37	2	post_office	16	2	1
community_centre	119	40	3	pub	135	21	3
dentist	104	35	3	restaurant	135	32	3
doctors	60	42	2	salon	141	53	3
embassy	75	24	2	shelter	90	0	2
fast_food	31	15	2	shop	43	21	2
grocery	20	10	1	spa	161	54	3
gym	100	22	3	stripclub	140	46	3
hookah_lounge	130	17	3	studio	60	0	2
ice_cream	23	7	1	veterinary	67	29	2
karaoke	188	15	3	vintage_modern_resale	38	32	2

Table 3: All amenities listed with their corresponding mean visiting duration (in minutes) and standard deviation, as collected from Google Places. The assigned category for cluster vectors is included.

Source ID ↕	Model ↕	Training-Error** ↕	Target ID ↕	Test-Error** ▲	Similarity ↕	Sim-Type ↕
5	dt	14.76	2	11.66	1.00	cosine
5	dt	14.76	8	14.12	0.99	cosine
5	dt	14.76	0	16.40	0.99	cosine
5	dt	14.76	7	22.53	0.98	cosine
5	dt	14.76	4	22.77	0.93	cosine
5	dt	14.76	9	22.93	0.98	cosine
5	dt	14.76	1	26.63	0.93	cosine
5	dt	14.76	3	27.93	0.85	cosine
5	dt	14.76	6	29.61	0.99	cosine
*Cluster Datapoints: 9.112						
**Error Type: Root Mean Square						

Figure 8: The pop-up table for the Leaflet application view of Figure 7.

Cluster ID	Data Points	Occupancy Points	Blocks per Timestamp
0	9879	120 320	12.2
1	12 387	203 728	16.4
2	8713	61 839	7.1
3	6134	22 371	3.6
4	9586	110 463	11.5
5	9112	87 316	9.6
6	10 244	118 096	11.5
7	9500	115 588	12.2
8	9051	112 245	12.4
9	8713	76 230	8.7
Average	9332	102 820	10.5

Table 4: Number of training points per model alongside the initial occupancy points within the containing blocks.

coefficients: the Pearson correlation coefficient and Spearman's rank correlation coefficient.

We have evaluated both cosine and EMD similarity values in configurations of 8 and respectively 16 clusters. Additionally, we varied the *merge distance* to see how the correlation behaves. The similarity values are hence calculated for 100m, 200m, and 400m merge distance respectively. In Table 8 the final results are shown. For each correlation measure, the percent of similarity values were calculated that correlated positively, in case of EMD distance, or negatively, for cosine similarity, with the estimation errors. The models taken were trained with gradient boosted trees.

We notice that the cosine similarity achieves better results than EMD for the same testing configuration, peaking at 100% negative correlation for 8 clusters and 100m merge distance. Its average Pearson coefficient is -0.55 , while the mean Spearman rank coefficient is -0.49 . EMD positively correlates the most for the same testing configuration, when the average Pearson coefficient is at 0.28 and Spearman's rank coefficient equals 0.23. There is a clear descending trend in correlations, as the merge distance increases. Also, the results for 8 clusters are superior to the ones when the city is split in 16 clusters.

C_{target}	$M(\mathcal{A} \setminus C_{target})$	$error_{total}$	$error_{cluster}$	$M(C_{source})$	C_{source}
0	xgb	18.20	16.10	xgb	6
1	dt	17.44	12.14	xgb	7
2	xgb	20.33	17.45	svm	5
3	xgb	17.59	13.35	xgb	1
4	dt	17.49	17.30	xgb	6
5	xgb	18.44	16.08	xgb	2
6	xgb	16.00	15.92	svm	0
7	dt	16.38	10.87	xgb	3

Table 5: Comparison between the estimations of the total models versus the ones of clustered models, in case the city is split into 8 clusters. Errors are expressed as RMSE.

C_{target}	$M(\mathcal{A} \setminus C_{target})$	$error_{total}$	$error_{cluster}$	$M(C_{source})$	C_{source}
0	mlp	19.53	13.40	xgb	15
1	xgb	17.82	17.18	xgb	7
2	xgb	20.79	17.97	svm	6
3	xgb	16.44	13.25	xgb	0
4	xgb	17.58	16.40	xgb	8
5	dt	17.59	10.93	xgb	0
6	xgb	18.44	16.29	xgb	2
7	xgb	16.73	16.38	xgb	9
8	xgb	13.57	14.14	xgb	10
9	xgb	16.90	15.95	xgb	10
10	xgb	16.25	16.45	xgb	7
11	mlp	21.73	14.95	xgb	11
12	xgb	20.33	15.42	xgb	0
13	dt	14.52	11.39	xgb	15
14	dt	22.93	18.73	svm	6
15	dt	20.63	13.33	xgb	0

Table 6: Comparison between the estimations of the total model versus the ones of clustered models, in case the city is split into 16 clusters. Errors are expressed as RMSE.

	dt	svm	mlp	xgb
8 clusters	24.6%	17.5%	12.3%	45.6%
16 clusters	14.6%	13.8%	13.8%	57.9%

Table 7: The proportion of best models by means of machine learning method.

6.10 Estimations for Clusters *without* Parking Data

We apply the models trained on SFpark data on clusters *without* parking data. The testing data records are composed of values equal to the averages of the respective data types in all clusters *with* parking data. This is the case for *parking price* and *parking capacity*. One piece of data that still needs to be provided so that the estimation is computed is the date and time. For convenience,

we choose the next day at the point when the user starts the model training. An example of the estimation is visualized Figure 9.

7 FUTURE WORK

To further investigate parking occupancy prediction given the assumptions in this work, there are several improvements or alternative approaches that can be realized.

- (1) **Use other parking settings.** In the present work, several pieces of data could not be integrated because of merging issues, i.e., the location unit did not coincide with the occupancy data's block. Traffic, events, weather, etc. could improve estimation results and hence the final estimations for clusters without parking data. Other sources for parking occupancy data can be found for the cities of Cologne [12], Zurich [24], Santa Monica [15]. In Germany, Deutsche Bahn provides an API to obtain data from parking around train stations [8]. Data pertaining to street occupancy is, however, hard to find. At the time of writing, open data portals mostly

C-WOUT ID 19 Timepoint: 2017-11-04 09:00:00 ▾
 Similarity [cosine](#) | [emd](#)

C-WTIH ID	Similarity	Sim-Type	Estimation	Interval	Intersection
6	1.00	cosine	45.27	[45.19 - 45.35]	[45.19 - 45.35]
7	1.00	cosine	45.15	[45.01 - 45.29]	[45.19 - 45.29]
8	0.99	cosine	57.73	[56.99 - 58.46]	empty
9	0.99	cosine	50.84	[49.84 - 51.83]	empty
5	0.99	cosine	66.71	[65.49 - 67.93]	empty
0	0.98	cosine	52.52	[50.58 - 54.46]	empty
1	0.98	cosine	36.32	[34.10 - 38.54]	empty
2	0.98	cosine	55.45	[53.03 - 57.87]	empty
4	0.97	cosine	58.11	[55.53 - 60.68]	empty
3	0.92	cosine	45.26	[37.51 - 53.01]	empty

Figure 9: The pop-up table of a cluster without data. Notice the drop-down list from which the time can be selected.

	8 clusters			
merge distance	cosine	rank_cosine	emd	rank_emd
100m	100%	100%	87.5%	75%
200m	75%	75%	75%	75%
400m	62.5%	50%	75%	75%

	16 clusters			
merge distance	cosine	rank_cosine	emd	rank_emd
100m	75%	75%	68.75%	62.5%
200m	75%	75%	56.3%	56.3%
400m	68.8%	68.8%	62.5%	62.5%

Table 8: Correlations between similarity values and model estimations errors for pairs of clusters *with* parking data (C_{source}, C_{target}). For cosine similarity, the proportions show the negative correlation, while for EMD, they express the positive correlation. The Pearson coefficient and Spearman’s rank coefficient were used.

provide the location of parking lots, parking meters, parking price and opening times, if applicable.

- (2) **Use more city data.** The parking profiles in the present work are relying on the public amenities from OpenStreetMap. OSM has great potential as a collaborative map service but it lacks many pieces of information that could be useful. Data such as opening hours would be interesting to include in the parking demand profile, which would then take into account the number of public amenities that are available at a certain point in time. Furthermore, the stay duration data collected for the present approach is currently limited. Adding more stay duration data may fine tune the emerging similarity values. Overall, more and finer city data, together with an appropriate representation and similarity function could

eventually improve the occupancy estimations for clusters without parking data.

- (3) **Integrate city data in machine learning models.** An alternative to building urban measures and similarity functions is to let machine learning figure out the similarities by itself. One can add the city data as further training information for clusters. The models are then applied on clusters without parking data and return occupancy estimations. The difference here is that models will be built for all clusters, including the ones without parking data. This also has the disadvantage of not using most of the parking data for training. The benefit of finding better similarities by leveraging unknown patterns in the city data may, however, outweigh this drawback.
- (4) **Apply semi-supervised machine learning.** Another relevant machine learning approach in this case is based on organizing the city areas as an undirected graph. The vertices represent the clusters with their respective occupancy data, while the edges between them are assigned similarity values. Initially, only a part of the vertices have the occupancy value known, i.e., the clusters with parking data, while the rest has undetermined occupancy, i.e., the clusters without parking data. At each step, the value for a vertex whose value is undetermined is being computed by considering the occupancies of the linked vertices and their corresponding similarity values.

8 CONCLUSION

In this work, we have presented an attempt at approximating street parking occupancy in cities. Under the assumption that parking data is lacking, in order to build scalable occupancy prediction systems, we proposed an alternative solution to the ones previously developed for this problem. We built parking demand profiles by using complementary city data, which localize various types of public

amenities and indicate the average visiting duration there. All data has been made available in an established RDF format, so that it can be easily reused. We merged the *parking data* with the *city data* by matching parking *location units* to *points of interest*, split the city into clustered areas, and built machine learning models for them. K-Means was used to cluster the location units, while four methods were employed to train models for the clusters: decision trees, support vector machines, multilayer perceptrons and extreme gradient boosting. Based on the city data, urban measures were built in the form of *cluster vectors* and *cluster Gaussians*, both of which took advantage of the mean visiting duration and its standard deviation. The vectors were part of the *cosine similarity* computation, while Gaussians contributed to the *earth mover's distance* calculation. The occupancy estimations for clusters *without* parking data were defined in terms of model estimations from clusters *with* parking data and the corresponding cluster similarity values. The estimations are expressed as intervals which extend the model prediction values by the magnitude of the similarity values.

As use case, we chose the SFpark project from San Francisco, which gathered parking data for more than 2 years starting in 2011 and now offers it for free usage. The city data was collected from OpenStreetMap as amenity information, and from Google Places as stay duration values. Both sources are open and free of charge. Over 30 types of public amenities were found in the San Francisco blocks, which corroborated with over 470 Google Places sources, lead to building the urban measures and similarity values.

The results confirmed that clustering the city into smaller areas yields better occupancy estimations than those of entire city area models. Following our tests, the best machine learning model turned out to be *extreme gradient boosting*. We used the clusters *with* parking data for the evaluation of the similarity values and calculated correlation coefficients between the similarity values and the estimation errors, using both absolute values and ranks. The best correlation were reached for the 100m merge distance for 8 clusters, averaging at -0.55 as Pearson Coefficient and -0.49 as Spearman's rank coefficient. In the same configuration, both *cosine similarity* and *EMD distance* reached their best results from all the test configurations. Overall, *cosine similarity* achieved better correlations than *emd*. Finally, the models for 8 clusters produced superior results over the models for 16 clusters.

REFERENCES

- [1] San Francisco Municipal Transportation Agency. 2011–2013. SFpark. (2011–2013). <http://sfpark.org>
- [2] San Francisco Municipal Transportation Agency. 2011–2013. SFpark - Open Data. (2011–2013). <http://sfpark.org/how-it-works/open-data-page/>
- [3] San Francisco Municipal Transportation Agency. 2011–2013. SFpark - Pilot Project Evaluation Summary. (2011–2013). http://sfpark.org/wp-content/uploads/2014/06/SFpark_Eval_Summary_2014.pdf
- [4] Felix Caicedo, Carola Blazquez, and Pablo Miranda. 2012. Prediction of parking space availability in real time. *Expert Systems with Applications* 39, 8 (2012), 7281–7290.
- [5] Murat Caliskan, Andreas Barthels, Bjorn Scheuermann, and Martin Mauve. 2007. Predicting parking lot occupancy in vehicular ad hoc networks. In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th. IEEE*, 277–281.
- [6] Xiao Chen. 2014. Parking occupancy prediction and pattern analysis. *Dept. Comput. Sci., Stanford Univ., Stanford, CA, USA, Tech. Rep. CS229-2014* (2014).
- [7] Zhirong Chen, Jianhong Cecilia Xia, and Buntoro Irawan. 2013. Development of fuzzy logic forecast models for location-based parking finding services. *Mathematical Problems in Engineering* 2013 (2013).
- [8] DB BahnPark GmbH. 2016. DB - Parkplätze API. (2016). <http://data.deutschebahn.com/dataset/api-parkplatz>
- [9] Reinhard Hossinger, Klaus Heimbuchner, and Tina Uhlmann. 2012. Development of a Real-time Model of the Utilisation of Short-term Parking Zones. In *19th ITS World Congress*.
- [10] Andrei Ionita. 2017. Extending Estimation of Parking Occupancy to Untracked City Areas using City Background Information. (4 Dec. 2017).
- [11] Andrew Koster, Allysson Oliveira, Orlando Volpato, Viviane Delvequio, and Fernando Koch. 2014. Recognition and recommendation of parking places. In *Ibero-American Conference on Artificial Intelligence*. Springer, 675–685.
- [12] Stadt Köln and DKAN. 2015. Öffene Daten Köln. (2015). https://www.offenedaten-koeln.de/dataset/taxonomy/term/52/field_tags/Transport%20und%20Verkehr-52?query=park&sorting=changed%7CDESC
- [13] Suhas Mathur, Tong Jin, Nikhil Kasturirangan, Janani Chandrasekaran, Wenzhi Xue, Marco Gruteser, and Wade Trappe. 2010. Parknet: drive-by sensing of roadside parking statistics. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 123–136.
- [14] Anandathirtha Nandugudi, Taeyeon Ki, Carl Nuessle, and Geoffrey Challen. 2014. Pocketparker: Pocketsourcing parking lot availability. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 963–973.
- [15] City of Santa Monica. 2014. Santa Monica - Open Data. (2014). <https://data.smgov.net/Transportation/Parking-Lot-Counts/ng8m-khuz>
- [16] University of Surrey (UK), Alexandra Institute (Denmark), Ericsson (Sweden), National University of Ireland/DERI (Ireland), University of Applied Sciences Osnabrück (Germany), Siemens (Austria, Romania), City of Aarhus (Denmark), City of Braşov (Romania), and Wright State University (USA). 2013–2016. City Pulse. (2013–2016). <http://www.ict-citypulse.eu/page/>
- [17] Tooraj Rajabioun, Brandon Foster, and Petros Ioannou. 2013. Intelligent parking assist. In *Control & Automation (MED), 2013 21st Mediterranean Conference on*. IEEE, 1156–1161.
- [18] Tooraj Rajabioun and Petros A Ioannou. 2015. On-street and off-street parking availability prediction using multivariate spatiotemporal models. *IEEE Transactions on Intelligent Transportation Systems* 16, 5 (2015), 2913–2924.
- [19] Felix Richter, Sergio Di Martino, and Dirk C Mattfeld. 2014. Temporal and spatial clustering for a parking prediction service. In *Tools with Artificial Intelligence (ICTAI), 2014 IEEE 26th International Conference on*. IEEE, 278–282.
- [20] Donald Shoup. 2001. Free parking or free markets. (2001). <https://www.accessmagazine.org/spring-2011/free-parking-free-markets/>
- [21] Matthew Nitch Smith. 2016. The number of cars worldwide is set to double by 2040. (2016). <https://www.weforum.org/agenda/2016/04/the-number-of-cars-worldwide-is-set-to-double-by-2040>
- [22] Tim Tiedemann, Thomas Vögele, Mario Michael Krell, Jan Hendrik Metzen, and Frank Kirchner. 2015. Concept of a Data Thread Based Parking Space Occupancy Prediction in a Berlin Pilot Region.. In *AAAI Workshop: AI for Transportation*.
- [23] Bo Xu, Ouri Wolfson, Jie Yang, Leon Stenneth, S Yu Philip, and Peter C Nelson. 2013. Real-time street parking availability estimation. In *Mobile Data Management (MDM), 2013 IEEE 14th International Conference on*, Vol. 1. IEEE, 16–25.
- [24] Stadt Zürich. 2015. Stadt Zürich - Open Data. (2015). <https://data.stadt-zuerich.ch/dataset/parkleitsystem>