Fleet Management in Free-Floating Bike Sharing Systems Using Predictive Modelling and Explorative Tools

Matthias Templ
ZHAW

Christoph Heitz
ZHAW

Abstract

For redistribution and operating bikes in a free-floating systems, two measures are of highest priority. First, the information about the expected number of rentals on a day is an important measure for service providers for management and service of their fleet. The estimation of the expected number of bookings is carried out with a simple model and a more complex model based on meterological information, as the number of loans depends strongly on the current and forecasted weather. Secondly, the knowledge of a service level violation in future on a fine spatial resolution is important for redistribution of bikes. With this information, the service provider can set reward zones where service level violations will occur in the near future. To forecast a service level violation on a fine geographical resolution the current distribution of bikes as well as the time and space information of past rentals has to be taken into account. A Markov Chain Model is formulated to integrate this information.

We develop a management tool that describes in an explorative way important information about past, present and predicted future counts on rentals in time and space. It integrates all estimation procedures. The management tool is running in the browser and continuously updates the information and predictions since the bike distribution over the observed area is in continous flow as well as new data are generated continuously.

Keywords: free-floating bike systems, spatio-temporal data, prediction, generalized additive models, markov chain, managment tool.

1. Introduction

For a service provider, redistribution costs represent a significant expense for operating a free-floating vehicle fleet. Heitz, Etschmann, Stöckle, Bachmann, and Templ (2019) and Lippoldt, Niels, and Bogenberger (2019) showed the effectiveness of incentivizing users via reward zones in the urban area. The aim is that service operators minimize the redistribution of bicycles (operator-based redistribution), by motivating users through a system of rewards to distribute the bicycles in the urban area in such a way that the service level is maximized (user-based redistribution). For setting such incentives, the current bike distribution as well as the predicted changes by picking and dropping of bikes has to be known (see Heitz *et al.* 2019). An important input for any incentive system is the knowledge of the spatial-temporal

distribution of future picks and drops, respectively. However, both Heitz *et al.* (2019) and Lippoldt *et al.* (2019) didn't show nor integrated forecasts in their model based on all available information - the current bike distribution and historical booking data.

For the task of redistribution, not the physical distribution of bicycles in a service area is the most important factor, but the availability of bikes in case of users requesting one. The goal of redistribution is to maximize satisfied demand and to avoid and service level violations due to unsatisfied demand. Therefore, the estimation of the future demand is of primary importance for any redistribution activity.

A demand model based on booking data is difficult to set up in free-floating systems due to different reaons:

a) even if months of historical booking data is available, these data are sparse because of the many degrees of freedom of an spatial-temporal distribution.

b) the picking data does not directly reflect the user demand, because users typically have to walk to pick up a bike. So, the real demand occurs somewhere else than where the actual pick happens and is recorded. Studies have shown that a typical acceptable walking distance is around 300-500 meters Singla, Santoni, Bartók, Mukerji, Meenen, and Krause (2015).

c) unsatisfied demand is not measured.

For dropping, similar issues arise. For example, the dropping place does not directly reflect the real destination of a trip, because bikes may only be parked in allowable areas (e.g. not in buildings or on the road).

For our model, we do not attempt to estimate the real demand (c) and assume that unobserved unsatisfied demand is spread homogeneus over the whole area. And we outline in the conclusion that unobserved demand is part of future research. However, we fully consider (a) and (b). We estimate a spatio-temporal picking rate and dropping rate, respectively. We assume that picking of a parked bike at (x, y, t) is done according to an inhomogeneous Poisson process the the estimated rate $\lambda(x, y, t)$, and that a picked bike is returned at a place given by the dropping distribution $\mu(x, y, t)$.

A simple demand model, as used, for example, in Reiss and Bogenberger (2016), only works if the data is not sparse, i.e. if only a very rough partition of an area in zones is applied. The authors claim that the demand is the average rentals in a time interval in a certain zone divided by the stock reported in the corresponding area. However, their model has no fine spatial and time resolution as we suggest and formulate in this contribution.

In the following sections we will see that for a service level violation estimate, the current distribution of bicycles (Section 2 and particularly Section 2.3) as well as smoothed rental statistics (Section 3 and especially Section 3.2) are important parts of the model. To improve short-term estimates, weather conditions can also be included in a model (Section 3.1). In order to show the important explorative statistics and model-based estimations in an explorative manner, a dashboard was developed which summarizes these results and makes them visible for the bicycle management (Section 5). Section 6 concludes and outlines further research.

2. Analysis

2.1. Descriptive figures and intraday distribution

The data is from the Swiss free-floating e-bike rental *Smide* (now *Bond Mobility*), based in Zurich. In the period between 2017-08-16 and 2018-03-22 a total of 63135 rentals from the free- floating e-bike sharing company *Bond Mobility* were booked in Zurich. *Bond Mobility* operates a fleet of speed pedelec (max. 45 km/h electrically-assisted pedal bicycles) which can

be booked using a smartphone app. The company services the fleet through in situ battery swaps and repairs/maintenance. The company does not re-position the vehicles manually. The vehicles are geographically balanced by its customers through regular use or through incentives provided in the form of free ride allowances. Incentives are awarded for positioning at charging stations or at locations where higher demand is expected. These locations can be determined algorithmically.

Especially the train stations in Zurich were most frequented, see Figure 3. As can be seen, most of the bikes were parked around the main train station of Zurich, at the ETH and the train station of Stadelhofen.

In the whole observational area, 385 bookings were made in average on a day (arithmetic mean, median equals 397) with a standard deviation of 130 (robust standard deviation through the median absolute deviation: 117). Guidon, Becker, Dediu, and Axhausen (2019), working with the same data from Zurich, argued that the main reason of the observed high standard deviation is because the system was growing during time. However, this effect is rather negatable. In contrast, we find that one of the most important reasons of this high standard deviation is in fact the weather condition. On days with bad weather condition, the number of bookings are smaller than on a sunny beautiful day. Figure 1a shows this effect only for precipitation intensity. Days with precipitation intensity below 0.1 result in an arithmetic mean of bookings of 17.5 per hour, while on days with high precipitation intensity the number of bookings was 13.9 (arithmetic mean), i.e. in average 26 % more bookings was made when the precipitation intensity was below 0.1. This is the reason for developing also a weather model (Section 3.1). Other effects are wind bearing and temperature, see Section 3.1.
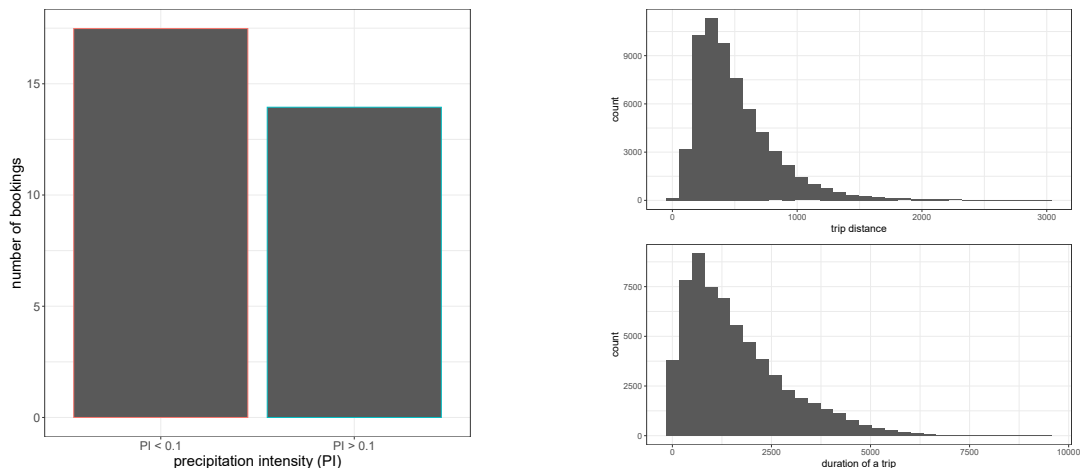
(a) Number of bookings on dry days and rainy days. (b) Trip distance (top) and trip duration (bottom).

Figure 1: Descriptive statistics on the number of bookings depending on the precipitation intensity, and the distribution of trip distance and trip duration.

Just to mention some interesting facts on mean trip distance and mean trip distance, even they have no further relavance in this contribution: The mean trip distance between picking and dropping of a bike was 1.198 km (95% confidence interval [1.192; 1.205], median = 1.08 [1070.475; 1091.532]). The mean duration of a trip was 8.94 minutes ([8.857; 9.041], median = 7.13 [7.083; 7.167]).

Figure 2 shows the temporal structure of the average total amount of trips per hour and weekday from 2017-08-16 to 2018-03-22. Figure 2 shows that a peak in the demand was reached 7:30 and 8:30 a.m (because of commuting), at noon (people using a bike to go for lunch) and between 5 and 6 p.m. (commuting home) during Monday till Friday. The supposed reasons for these peaks are supported by the fact that the pickings occur nearby train and other public transport stations (cf. Figure 3). The bookings on different weekdays behave very similar to each other. However, the bookings at weekend differs substantially in comparison

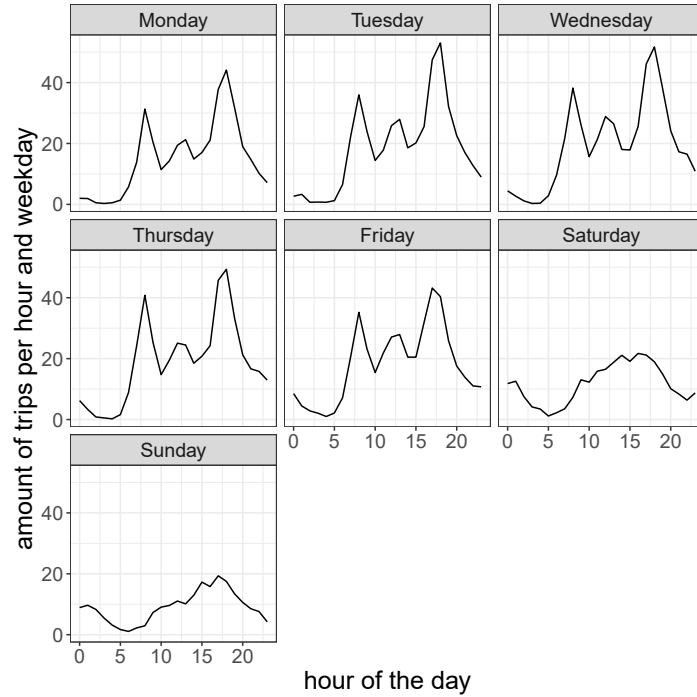to weekdays. Saturday and Sunday have similar frequencies.

Figure 2: Counts on bookings per hour for each weekday in the whole observational area.

Figure 3 shows the spatial distribution of drops. The upper left graphics shows the exact destinations of drops, the graphics in the upper right shows the summary on counts. This graphics are integrated in the management tool and zooming is possible. Moreover one can zoom until the location of drop of each individual bike is visible. The graphics in the lower left shows a two-dimensional kernel density estimate. Thus the density on drops is better visible using this map. The graphics to the lower right shows is zoomed to the area of Stadelhofen and drops (red points) are visible as well as the two-dimensional kernel density estimates. The drops mainly happen around the train station of Stadelhofen.

2.2. Dropping and picking rates

For controlling redistribution activities, it is essential to predict the spatial and time-dependent demand for bikes over the next few hours. In addition, the dropping rate must also be estimated as precisely as possible in order to forecast the change of the bike distribution. From these two pieces of information it is possible to estimate how the bike fleet and service level at a given location would change over the next few hours without external intervention. For our modelling we assume that for each location (x, y) in an area B and each time (t) an unknown, estimable dropping rate $\lambda_{drop}(x, y, t)$ and picking rate $\lambda_{pick}(x, y, t)$ exists. We assume that the actual process is an inhomogeneous Poisson process, i.e. that in every small time interval dt the probability of average intensity for a drop/pick in the interval $[t, t + dt]$ in area B is given by $m(t) = \int_B \lambda(x, y, t)dt$.

One task is to estimate this time- and location-changing rate from existing (historical) picking and dropping data. The number of drops and picks of *Bond Mobility*'s bicycles in temporally and geographically finely resolved areas is sparse, i.e. the counts on drops and/or picks are very small and mostly zero in each cell, even if several months of booking data is available (cf. Figure 4 and discussions related to Figure 4 below). Even in relatively well frequented areas
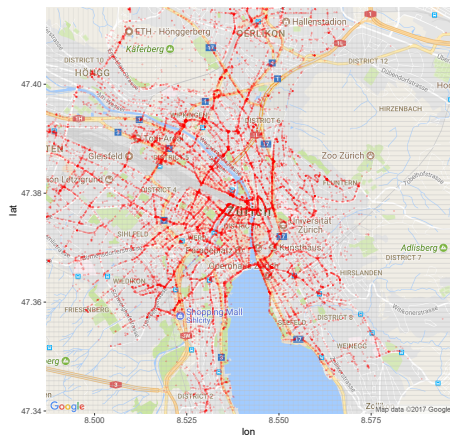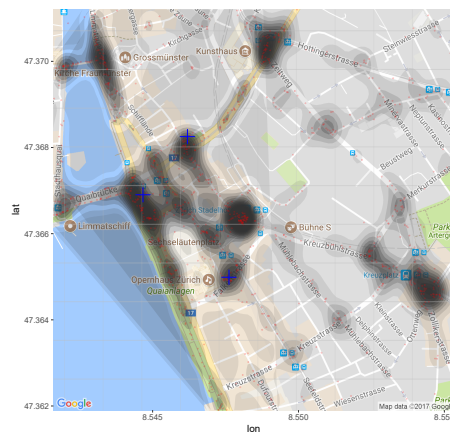
(a) Droppings in Zurich.

(b) Count density of drops.
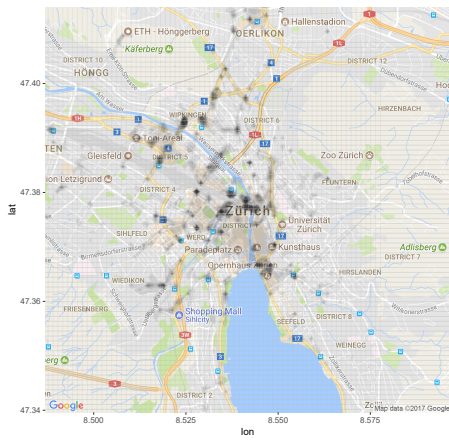
(c) Two-dimensional kernel density estimates.

(d) Drops and two-dimensional kernel density estimates around Stadelhofen (Zurich).

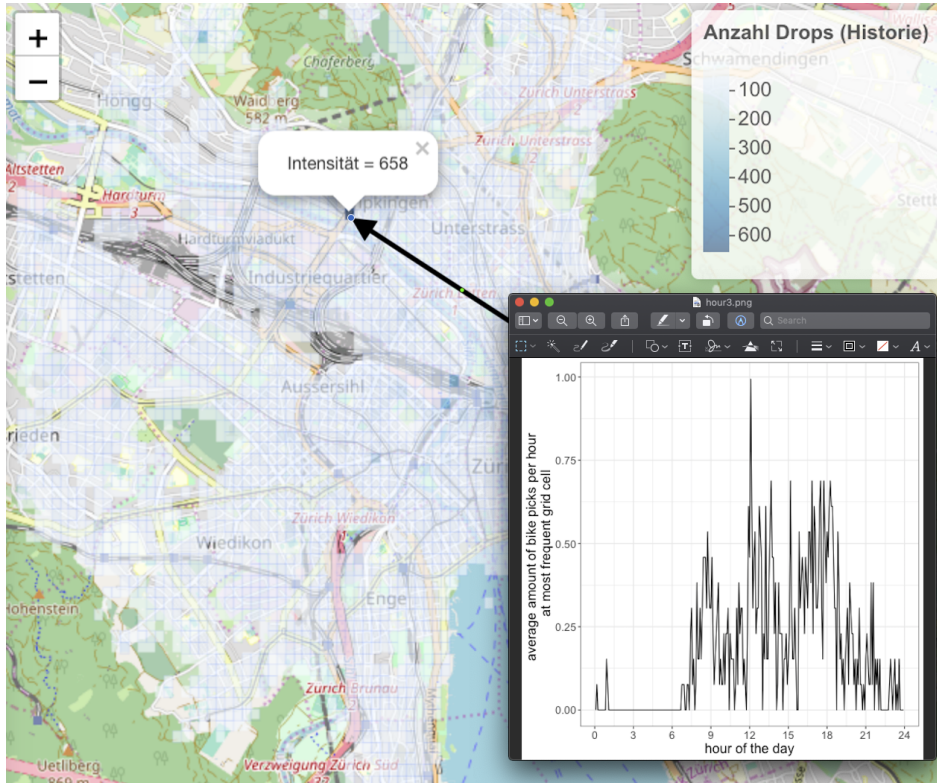Figure 3: Dropping of *Bond mobility* bikes.

Figure 4: Intensity of drops. The average rate of bookings per hour of one station (the most frequent one) is shown in the lower right.

one usually has to wait several hours until a pick or drop occurs. For the most frequent $100m \times 100m$ area of Zurich, the average amount of bike picks per hour is shown in Figure 4. The sum of all bookings in this cell over all observed days (see Section 3.1) is 658 bikes. About every 3-4 hours a bike is picked at this location. Note that this is the former place of the headquarter of *Bond Mobility*. The resulting small mean intensity makes it difficult to apply standard methods from the literature, e.g. modelling with a Cox model for inhomogeneous Poisson processes (see e.g. Diggle, Rowlingson and Su, 2005) and also agent-based mico-simulations (see, e.g., Ciari, Balac, and Axhausen 2016) or modelling using spatio-temporal point-process patterns (González, Rodríguez-Cortés, Cronie, and Mateu 2016; Baddeley, Rubak, and Turner 2015) cannot be performed or can only be performed with great difficulty.

Our model is based on a division of the observation area B as well as time of the day in a grid (see Section 3.2), predicting the rate in each (x,y,t)-cell. For the estimation, we allow different observation periods for the individual components of the model described below. Additional external factors/variables for the description of drops or picks as used in Müller, de Almeida Correia, and Bogenberger (2017); Guidon *et al.* (2019), e.g. data on inhabitants (age distribution, gender distribution, household composition) and areas (purchasing power, number of firms, hotels, services, rents) can be used to explain spatially distributed picks or drops. However, they are hardly useful for spatio-temporal forecasting, since these variables usually do not change during the observation period. More interesting is the inclusion of weather conditions, see Section 3.1. The size of the raster cells is chosen to be 100×100 metres. This represents a very fine grid and is in line with what was considered in Heitz *et al.* (2019), but a fine grid also implies that many grid cells contain no observations or only very few counts. These disadvantages can be eliminated with a well defined smoothing in time and space. Doing so, the advantages of making statements on a very fine grid remain.

Smoothed averaged dropping and picking rates per hour are shown in Figure 5 (see also Section 3.2). It can be seen that these rates are very small and that on most location it needs more than 10 hours for the next booking. Around the main train stations and also at the

ETH Zurich the rates are higher than in other areas. The graphics to the right of Figure 5 shows the total amount of bookings in the past during an hour at a given time of the day. This area is already zoomed and shows the area around the main train station of Zuerich down to the Zurich lake. At the ETH Zurich we see an intensity of 429 bookings.
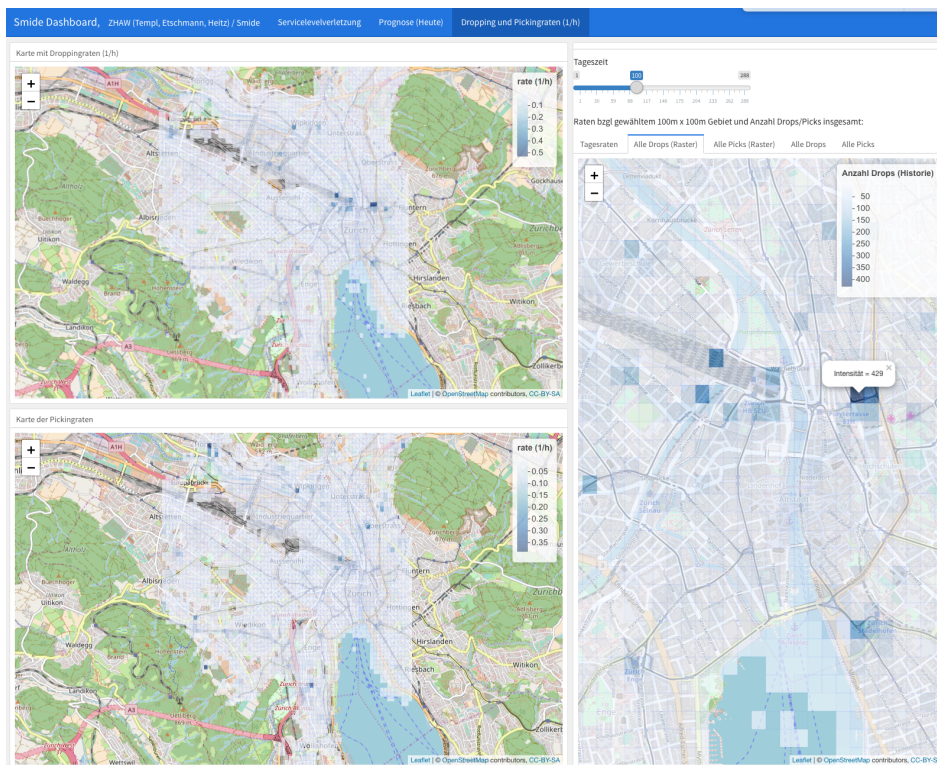

Figure 5: Rate of drops (upper left) and picks (lower left) per hour and the average intensity around the main train station in Zurich at a given time of the day.

2.3. Current bike distribution

The current bike distribution is read out of the *Bond Mobility* fleet operation system via an API, the data is processed and integrated on-the-fly. A result can be seen in Figure 6, which shows the current bike distribution that is queried from the server as soon as the app is started. We can see that most of the bikes, namely 40, are located around the main train station in Zurich. In this map, it is possible to zoom in, until all bikes can be seen individually (like the blue dot in the upper third of the graphic, which represents the exact location of a single bike).

The current bike locations are then needed in Section 4 to estimate the service level violation time on fine spatial and temporal resolution.

3. Three-component model

For estimating the dropping or picking rate, we suggest a decomposition of the spatial-temporal rate into three components:

$$\lambda(x, y, t) = \alpha I(x, y) g(t|x, y) \tag{1}$$

with $x, y \in B$ and time t. In the following, we only discuss the case of droppings, but the procedure is identical for the estimation of the picking rate.
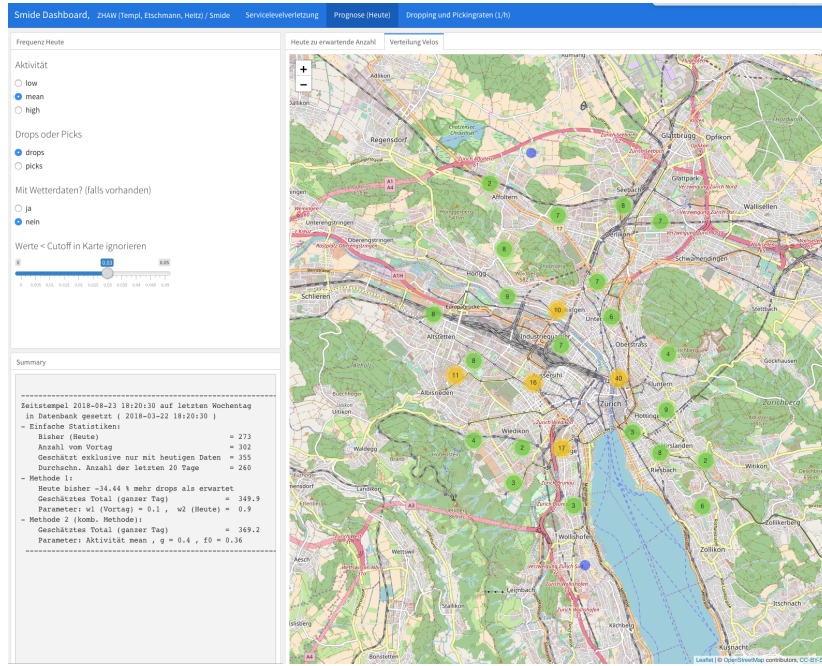
Figure 6: Current bike distribution of bikes from *Bond Mobility* in Zurich (accessed Feb 16, 2020, 17:12).

The parameter α denotes the total traffic intensity, $I(x, y)$ is the time-averaged spatial distribution of drops, and $g(t|x, y)$ is the intra-day pattern of a given location.

We assume that $I(x, y)$ and $g(t|x, y)$ do not depend on weather conditions, and that differences of the usage pattern due to external factors only influence the total traffic intensity α. This reduces the degrees of freedom substantially and leads to more stable estimators.

Meaning of α (number of picks or drops of bikes): α is the total number of drops during a full day. This parameter may vary from day to day, depending, e.g., on the weather conditions.

Meaning of $I(x, y)$: $I(x, y)$ indicates the relative proportion of drops that happen in cell (x, y). For example, $I(x, y) = 0.001$ means that 0.1% of all drops happen in cell (x, y). $I(x, y)$ is a dimensionless quantity which is normalized, so that $\sum_{x,y \in B} I(x, y) = 1$. Due to the normalization, $I(x, y)$ can also be interpreted as the probability that a drop will land in cell (x, y). Note that both α and $I(x, y)$ are quantities that correspond to aggregated numbers of drops: α corresponds to the number of drops for a full day, and $\alpha \cdot I(x, y)$ to the number of drops per day in cell (x, y).

Meaning of $g(t|x, y)$: $g(t|x, y)$ is the continuous temporal distribution of drops over the day at the point (x, y). Its dimension is the reciprocal of a time unit, e.g. h^{-1}, and it is normalized by $\int_0^{24h} g(t|x, y)dt = 1$, $\forall (x, y) \in B$. $g(t|x, y)dt$ can be interpreted as the probability that a drop in cell (x, y) will occur in the time interval $[t, t + dt]$. The intraday distribution may depend on the location.

This results in the following important partial rates: Total number $N(x, y)$ of drops at the location over a full day $N(x, y) = \int_0^{24h} \lambda(x, y, t)dt = \alpha I(x, y)$

The total rate of drops at time of day t on a day is $\lambda(t) = \sum_B \lambda(x, y, t)$

Remarks: Any time- and location-dependent density can be represented in this way. The decomposition is therefore complete. The dropping rate $\lambda(x, y, t)$ at a given location (x, y) is

obtained by multiplying all three terms in equation above, in the correct units. Its dimension is "number of drops per time unit", where $g(t|x, y)$ is the term defining for the time unit.

Our decomposition is motivated by the possibility of using different time scales for the different components. For example, the spatial distribution $I(x, y)$ can be estimated with aggregated data of many days, ignoring the intraday structure and the traffic variability due to external conditions such as weather and temperature. Similarly, $g(t|x, y)$ can be estimated from aggregations of drops in a specific cell over many days.

In contrast, the total traffic intensity α (see Equation 1) can be estimated with a shorter observation period by aggregating over the whole region B. This allows the maximum information to be taken into account in the data, trends in the progressions to be depicted realistically and the sample to be kept as large as possible, depending on what is estimated.

3.1. Predicted intensity of the current day.

Three models are outlined to estimate the total amount of bookings (α) on a given day. The first two models do not need meterological information and they are thus easily applicable in practice. A better but more complex model (the weather model) integrate the meteorological information for prediction.

Except for days with different weather conditions, we can state that the normalized intensity of use for consecutive days is strongly correlated. This means that the previous day should be a good predictor of the current intensity of use (corrected with the weekday effect).

The first estimator for the number of picks of a given day d at time t_0 is given by

$$\hat{\alpha}_{d;\,t>t_0} = \alpha_{d;\,t<t_0} + w \cdot s \cdot \frac{c_{t_0}}{c_{t_0}+1} + (1-w) \cdot \alpha_{d;\,t<t_0} \cdot c_{t_0} \quad , \tag{2}$$

with $0 < w < 1$ a weight specified below. s is the average number of picks of the weekday estimated from historical data. For example, for the estimation of α on a day d, the mean of the drops on all previous days of the same weekday is used ($d-7, d-14, d-21, ...$). c_{t_0} defines a scaling factor between the average number of picks before time t_0 and after t_0 estimated from historical data on the given weekday ($d-7, d-14, d-21, ...$). $\alpha_{d;\,t<t_0}$ represents the observed number of picks from midnight until t_0 on the given day d. Thus, the prediction of the total traffic intensity is done on the basis of the drops already observed at day d until t_0, combined with the average intensity of similar day in the past.

It follows that the later in the day, the worse is the use of the average daily rate. The weighting constant w has thus to be selected. w is fixed to some constants. Before 8 a.m. the average of the historical data for the daily rate is used, thus w would be 1 in this case. From 9 o'clock on a mixture of the historical data and the extrapolated picks is used, whereby the historical data receives a weight of 0.4. For this choice, forecast error of 20-25% for the future picking rate can be achieved. The results of this approach is shown in the app (see Figure 6, lower left, method 2).

The second estimator for the number of picks of a given day is given by

$$\hat{\alpha}_{d;\,t>t_0} = w_1 \alpha_{d^*-1} + w_2 \cdot \bar{\alpha}_{(d-1,d-2,...,d-20)} \frac{\alpha_{d;\,t<t_0}}{\alpha_{(d-1,d-2,...,d-20);\,t<t_0}} \quad , \tag{3}$$

with $w_1 = 1 - t_0/t$ and $w_2 = 1 - w_1$, meaning a linear weight depending on the time of the day (with t indicating a whole day). Therefore the second choice (see Figure 6, lower left, method 1) gives linear weights depending on the time of the day. The first term in Equation 3 represents the weighted number of drops from the previous day. The star in d^* symbolizes that the previous day of Monday is the last Friday (must be a weekday) and the previous day of Sunday is the Saturday (because both are weekend days). It ensures that the previous day is thus the previous day with restriction to be again a weekday or weekend day respectively. The right side of Equation 3 expresses the average of previous 20 working (weekend) days

multiplied by the observed number of drops on day d until time t_0 divided by the average of the previous 20 working (weekend) days until time t_0 on a day. For example, on 18:20 most of the day is already passed, thus weight w_2 is much larger than w_1 in this case. The results of this approach is shown in the app as well (see Figure 6, lower left, method 1).

Estimation of α using meterological information: As all redistribution activities depend on short-term (typically with the next few hours) predictions of picking and dropping, the prediction model has to take into account the strong effects of external conditions such as temperature or precipitation. We assume that these factors influence the total traffic intensity α, while $I(x,y)$ and $g(t|x,y)$ do not strongly depend on them.

The following model is used to predict α, the total number of bookings of a day (e.g. 12. February 2020) at any time of the day (e.g. on 10:28).

A prior attempt to predict this quantity, using the same data set, was published by Guidon *et al.* (2019). The authors used a negative binomial least squares regression model without interaction effects resulting in an R^2 of 0.1 with the meterological information as predictors. It is not hard to see that all model assumptions must be violated using this model. The reason is that the intraday distribution of rentals is highly non-linear and may not be modelled using a linear model.

We formulated a generalized linear model (gam) on log counts and smoothing temperature, hour of the day, wind bearing and precipitation intensity with a thin plate b-splines smoother (Wood 2003, 2017) and integrating the weekday and seasonality (by month) in the model. This resulted in a R^2 of 0.7 and all diagnostic plots shows that the model assumptions are fulfilled.

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
(Intercept)	0.8471	0.2257	3.7532	0.0002
weekdayTRUE	0.3694	0.0321	11.4941	< 0.0001
month	0.1420	0.0239	5.9314	< 0.0001
B. smooth terms	edf	Ref.df	F-value	p-value
s(apparentTemperature)	8.3235	8.8739	9.8792	< 0.0001
s(hour)	8.9722	8.9997	252.4318	< 0.0001
s(windBearing)	6.9263	7.9983	2.3689	0.0153
s(precipIntensity)	8.2474	8.8250	3.4265	0.0003

Table 1: Results from the generalized additive model describing log counts on drops on all 100 times 100 raster cells in a hourly time resolution. These model coefficients are used for predicting the expected daily counts. The weather data for the Zurich area was obtained from DarkSky (https://darksky.net/).

Temperature, hour of the day, wind bearing and precipitation intensity had a highly significant effect, but also the weekend and seasonality (by month), see Table 1.

We use this model not only to describe and interpret the number of bookings as in Guidon *et al.* (2019), but also to predict the expected number of remaining bookings on a given day at a given time t_0 (i.e. the number of picks /drops that are expected for $t > t_0$, i.e. for the rest of the day), based on (1) the number of already observed total bookings until time t_0 of the current day and (2) predicted hourly bookings for the rest of the day using the weather forecast data and other predictors like hour, seasonality and weekday/weekend. The forumula can be summarized as

$$\hat{\alpha}_{d;\,t>t_0} = \alpha_{d;\,t<t_0} + \sum_{i \in \{1,2,\dots,24\};\,i>t_0} \hat{\alpha}_{d;\,i} \quad , \tag{4}$$

with $\hat{\alpha}_{d;\,i}$ the predicted total hourly number of drops on hour i to $i+1$ on day d using the model given in Table 1.

The weather model is thus a more elaborate model as the previous ones (Equations 2 and 3), but the drawbacks are that it needs a paid service from a meteorological service to fetch predicted hourly weather data. An advantage of this model using weather data is also that prediction intervals and not only point estimates can be reported. Naturally the prediction intervals become smaller as the time of the day increases (since more real observations on bookings are available).

3.2. Spatio-temporal smoothing

In this section, we describe the estimation of the remaining components of the three-component model, i.e. $I(x, y)$ and $g(t|x, y)$. It is important to determine how the drops/picks are distributed throughout the day. However, as already mentioned the data are sparse and smoothing is thus necessary.

For the following considerations, the observation area B is discretized in a grid of the size $n \times p$ in cells of equal size (e.g. 100m × 100m).

Our assumption is that the counts are spatially and temporally smooth. Small violations due to parking space (not everywhere it is allowed and possitble to drop a bike) does not play much of a role, because we smooth the drops (and picks) anyhow on $\pm 300m$ due to the walking of customers. We assume that the intraday change of counts does not change too much if we smooth over one or two hours. We also assume that the intraday variation does not change too much if we smooth over a range of $\pm 300m$ in all directions.

Another assumption is that the intraday variation is the same over all days of the week. Saturday/Sunday differs in daily variations (see also Figure 2). So we calculate two different intra-day cycles for each location (x, y).

Three-dimensional kernel smoothing The estimation method for weekday (Mon-Fri) is as follows, exemplarily noted here for dropping. We note that this procedure works the same for picking.

- Take all drop events on the weekdays Mon-Fri and assign them to the corresponding cell (x, y).

- Use only the time of day of the drops (e.g. 8:23)

- Divide the time within one day (24 h) into constant time intervals j of 5 min, thus $j = 1, ..., 288$. Note that also other lengths of time intervals can be chosen.

- Assign each drop to the corresponding time interval j.

- Result: Number of drops in each element (x, y, j). Note that often we rather work with relative number of drops, i.e. dividing with the number of drops in a day. This is also the case when estimating $I(x, y)g(t|x, y)$. For various graphics we show smoothed absolute numbers, e.g. in Figure 4.

- For each time interval j, perform a spatial kernel density estimate with a Gauss kernel of width 2 (i.e. spatial smoothing in the range of about $300m$).

- Then, for each position (x, y), perform a temporal kernel density estimation with a Gauss kernel of width 8 (i.e. temporal smoothing in the range of about $\pm 8 \cdot 5min$).

Except in peripheral areas, where the drops were very sporadic, this estimate provides a reasonable intraday variation, especially at the hotspots this works very well.

4. Time to service level violation

The most useful estimator is the time to service level violation for any location in the observational area B. Knowing this time to violation of the service at each place, the bike fleet can be efficiently redistributed, either by operator redistribution activities, or by proper user incentivation (Heitz *et al.* 2019; Heitz, Blume, Scherrer, Stöckle, and Bachmann 2020).

The time to a service level violation is based on the last state of the imported data and the estimated (smoothed) dropping and picking rates (see Section 3.2) from the last important state of the data import, as well as the current bike distribution (see Section 2.3).

For all control approaches for bike redistribution it is essential to know the local and time-resolved demand for bikes and to be able to predict it over the next few hours. Smoothed dropping and picking rates (per hour) are therefore determined and displayed based on a three-dimensional kernel density estimate (time, latitude, longitude) of the drops and picks with normal distribution kernel (see Section 3.2).

Given the booking rates per hour over d days and the current bike distribution, the expectation value for the time to the next service violation is estimated using a Markov chain model.

First the current bike distribution is fetched (see Section 2.3). Here, the exact position of $n \in \{1, ..., N\}$ bikes are imported, with N the total number of bikes from the service level provider. n is typically slightly smaller than N, because not all bikes are always in the field. We divided the whole service area in a raster of 100m, which results in $71 \times 101 = 7171$ cells. In addition, we divide the day in 5 minutes intervals leading to 288 time intervals for one day and do a prediction of service level violation for a future period of d days, for example $d = 5$.

We assume that a user would walk 300 metres for a bike at maximum, thus an available bike in one location (cell) will be available for a range of $\pm$ 300 metres. This is considered in the whole analysis.

For each of the 7171 cells (locations), we estimate the probabilities of picks given no bikes are available in the cell and neighboring cells (overbooking, service level violation), the probabilities of having no bike available (and no overbooking), the probabilities of having one bike in a cell or neighboring cell (within 300 metres), the probabilities of having two bikes in a cell in a location, etc., until the probabilities of having n bikes in a cell in a location. This information is stored in a matrix of probabilities $\mathbf{P}^{(7171 \times (n+2))} = (\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_{n+2})^T$.

The algorithm starts at time t_1 and ends in time t_{d*288} and updates another matrix holding the estimations on service level violation in every state of the process. Thus, in our calculation procedure we estimate these probabilities for every five minutes from the present time up to d days. For $d = 5$ days, this results in $288 \cdot 5 = 1440$ columns of $\mathbf{M} = (\mathbf{m}_1, \mathbf{m}_2, ..., \mathbf{m}_{288 \cdot d})^T$. In other words, the first column of $\mathbf{M}$ holds the probabilities of service level violation in the first five minutes, the second column of $\mathbf{M}$ the probabilities of service level violation within the first 10 minutes, etc.

More precisely, the probabilities of a service level violation is calculated iteratively for all times $t_2, ..., t_{d \cdot 288}$ by updating $\mathbf{P}$ with

$$
\begin{aligned}
\mathbf{P} &= \mathbf{P}(1 - droprate_{t-1} - pickrate_{t-1}) \\
&\quad + \mathbf{P}^{(*2)} pickrate_{t-1} \\
&\quad + \mathbf{P}^{(*p)} droprate_{t-1} \\
\mathbf{p}_1 &= \mathbf{p}_1 + \mathbf{p}_2 pickrate_{t-1} \\
\mathbf{p}_2 &= \mathbf{p}_2(1 - droprate_{t-1} - pickrate_{i-1}) + \mathbf{p}_3 \cdot pickrate_{t-1} \\
\mathbf{p}_p &= \mathbf{p}_p(1 - pickrate_{t-1} + \mathbf{p}_{p-1} droprate_{t-1}) \\
\mathbf{m}_t &= \mathbf{p}_2 pickrate_{t-1}
\end{aligned}
\tag{5}
$$

whereby $\mathbf{P}^{(*2)}$ is the reordered variables of matrix $\mathbf{P}$ with column order $(2, 3, ..., d \cdot 288, 1)$ and $\mathbf{P}^{(*d \cdot 288)}$ the reordered variables with column order $(d \cdot 288, 1, 2, 3, ..., d \cdot 288 - 1)$. The first line of Equation 5 represents the probability at time t of having zero bikes reachable (within 300 metres of a location) under the contraint that in the last 5 minutes a bike has been dropped

or picked. The second line represents the probability that two bikes was reachable (within 300 metres) and one was picked. The third line adds the probability of having zero bikes available but one is dropped. For the situations of having already an overbooking ($\mathbf{p}_1$) we overwrite the estimated probabilities, because in this special case, overbooking is expressed as the probability of having an overbooking plus the probability of having a pick but no bikes. Also the case $\mathbf{p}_1$ (no bikes available) must be treated differently as the rest. Here it is the probability of having no bike available times the probability of the complementary event of a bike being dropped or picked plus the probability that exactly one bike is available and a pick is happening. Note that the drop and pickrates are estimated with the three-dimensional kernel estimation procedure decribed in Section 3.2.

In theory, the algorithm can iterate over infinite days. However, at one point for sure a service level violation happens and, in addition, the most interesting case for a service provider is to know the service level violations in the next 6 to 48 hours. This is usually enough time to redistribute bikes.

Results (accessed on Feb 16, 2020, 17:12) are shown in Figure 7. Since the current bike distribution (Figure 2.3) shows a number of bikes in the area of the main train station and ETH Zurich, there is no service level violation in the next hours in these regions. Left to Riesbach, for example, there were no bikes available (Figure 2.3) and given the past smoothed booking rates, we expect a service level violation soon. As another example, there was only four bikes available between in the area of Zurichberg (Figure 2.3), but given our Markov model using smoothed the historical booking rates, we do not expect any service violations in the next 60 hours.
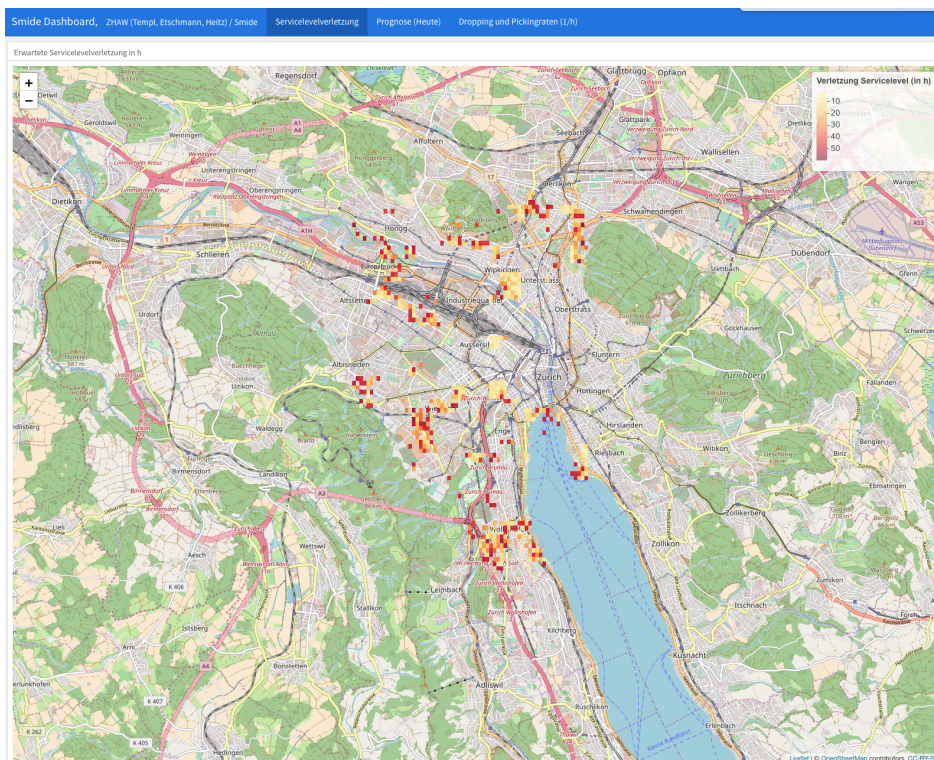

Figure 7: Service level violation in the next 0 to 60 hours estimated by the Markov model given the current bike distribution and smoothed historical picking and dropping rates. The darker the colour the longer it needs for a service level violation.

5. Monitoring

For supporting the operator in planning its redistribution activities as as well as setting reward zones for user incentivization, we developed a novel interactive map, displaying the following information:

- The graphics in this article represents static screenshots of interactive graphics obtained from the app and are used to capture what can be viewed in the app. The app can be operated interactively in the browser and all graphics are highly interactive.

- All maps can be zoomed and moved in all directions.

- All maps are reactive. By mouse clicks on points shown on the map, additional information showing some estimates and current figures pops up.

- The drops and picks of the past as well as the current bike distribution are processed internally.

Output are therefore different reactive graphics which change partially when the user input changes and where pop-up info is displayed when pointing with the mouse on the map.

We partioned the information in different tabs.

5.1. Service level violation

This tab shows the expected values of the time to the next service level violation per cell, see a screenshot in Figure 7. These expectation values indicate how long it will take on average until a demand cannot be satisfied in this cell. The lower this value is, the more urgent a bike is needed in this area. Therefore, reward zones can be placed in cells that have a very low expected value in order to incentivize users to drop their bike in this region. Alternatively, the operator might become active and position a bike in this region.

For the calculation of the expected values, the current bike distribution (regarding the last update of the database, see Section 2.3) and the estimated dropping and picking rates (see Section 3.2) are used. With the rates over 5 days and the bike distribution the mean values are calculated using the Markov chain. In the browser, one can zoom and mouse pop-ups show additional information (exact estimated time to service level violation, droping and picking rates, name of the area, etc.).

5.2. Daily forecast and current bike distribution

Figure 8 shows the forecast of the number of drops or picks over the service area.

The user may manually chose between whether he expects high activity, medium activity or high activity, and can toggle the display between picking and dropping. For low or high activity the 20% and 80% quantiles of the distribution of daily picks or drops are used.

To get a picture focused on higher values, small values can be truncated in the graph. This has the advantage to focus on high values in the map.

In the displayed summary simple statistics regarding the drops or picks are shown. For example, it is reported for a particular day that until 18:20 the number of drops are 273. Also the number of drops from the day before are shown as well as the estimated number of drops using only the extrapolation until the end of the day. Also the average number of drops (or picks) of the last 260 days are shown. Method 1 corresponds to a weighted estimate with automatic linear weighting - the later the time of the day, the higher the weighting to today. Since it is already 18:20, the weighting is higher (0.9) for the 273 bikes until 18:20 and less weight (0.1) is given to the weight corresponding to the previous day (302 bookings)
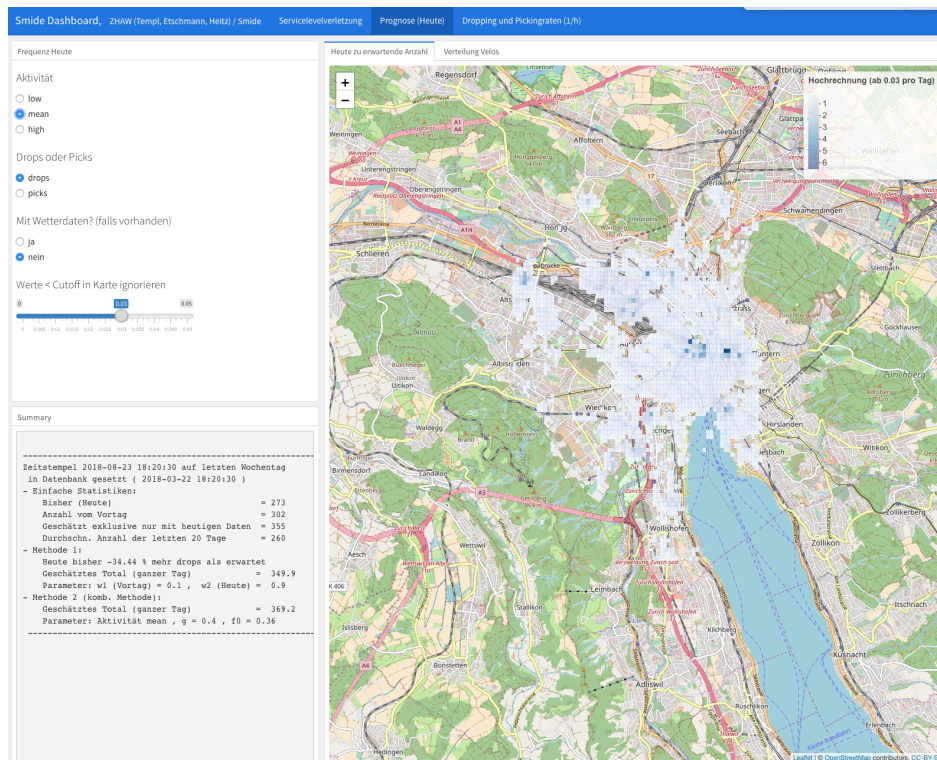
Figure 8: Number of drops (or picks) per day in each grid cell (right graphics) and daily forecasts for the absolute number of drops (or picks) in a day (upper left).

for the estimation of the daily bookings. Method 2 contains the method after fixed weighting according to the the 3-component model from Section 3.1. Method 2 expects for this given day 369.2 bookings. In case meterological data from a from a paid service are available for each hour of all days in the past and predictions from the meterological service for the next hours, a third method pops up showing the estimated number of bookings from the weather model.

Again, one can zoom in and out and get popup information (browser only).

The current distribution of the bikes can be retrieved in the tab *Distribution* of Bikes, see also Figure 6. The displayed numbers are broken down more finely (up to the exact location of each individual Bicycle) when zooming in.

5.3. Daily forecast

The dropping and picking rates (per hour) are determined and displayed based on a three-dimensional kernel density estimate (time, latitude, longitude) of the drops and picks with normal distribution core, see Section 3.2. A screenshot of the corresponding results are shown in Figure 5. Note that the right graphics shows the estimates on a given time of the day, and the time of the day can be changed with a slider with stepsize of 5 minutes. The rates shown in this figure are also input for the service level violation estimation in the first tab of the app.

It can again be zoomed as desired and popup info is available in the browser. The time of day can be changed interactively.

In the graphic on the right you can also see information about the intensity of past total drops and total picks including popup information, already shown here as an example.

The intensity is also indicated in further interactive graphics, see the screenshots represented in Figure 3.

6. Conclusion

The existing literature on spatial-temporal usage patterns of bike sharing systems mainly focuses on the analysis of historical bookings using descriptive statistics or models. For controlling and optimizing any redistribution activities, the more essential question is where a service violation is expected in the near future given the current distribution of bicycles and given the temporal and spatial foecast of both demand and bike distribution.

In other words, for redistribution and operating a bike fleet in a free-floating system, one must be aware of short term forecasts on daily bookings and of forecasts of spatial and time-resolved service level violations that are expected to occur if no redistribution activities are triggered. In addition, a fleet manager also needs a management tool (dashboard), which makes these predictions visible in maps and easy-to-fetch figures. Based on this information, the most appropriate redistribution activities can be set either automatically or by expert rating.

We described a management tool (dashboard) which is running in a browser, continuously updating the information and predictions. This tool thus provides a decision tool for the fleet managers that allows them to optimally redistribute bicycles in orderto avoid service level violations.

Further research may be addressed due to unobserved demand issues. A likelihood-based approach can be used to model such unobserved demand, e.g. to specify the likelihood of having an unobserved demand given there is no bike available in a 300 metres range. The parameters of the likelihood functions must then be estimated accordingly based on observed picking rates.

Note that the described Markov chain model (see Section 4) is currently used by the company roll2go (https://www.roll2go.ch) for automatically setting reward zones based on the presented service level violation estimation procedure. The R package mobi2 (Templ 2019) (available upon request) includes all estimation procedures presented in this article and it also includes the introduced shiny app (the dashboard, managament tool).

References

Baddeley A, Rubak E, Turner R (2015). *Spatial Point Patterns: Methodology and Applications with R*. Chapman & Hall/CRC Interdisciplinary Statistics. CRC Press. ISBN 9781482210217.

Ciari F, Balac M, Axhausen K (2016). "Modeling Carsharing with the Agent-based Simulation MATSim: State of the Art, Applications and Future Developments." *Transportation Research Record*, **2564**(1), 14–20. doi:10.3141/2564-02.

González JA, Rodríguez-Cortés FJ, Cronie O, Mateu J (2016). "Spatio-temporal Point Process Statistics: A Review." *Spatial Statistics*, **18**, 505 – 544. ISSN 2211-6753. doi:10.1016/j.spasta.2016.10.002.

Guidon S, Becker H, Dediu H, Axhausen K (2019). "Electric Bicycle-Sharing: A New Competitor in the Urban Transportation Market? An Empirical Analysis of Transaction Data." *Transportation Research Record*, **2673**(4), 15–26. doi:10.3929/ethz-b-000324494.

Heitz C, Blume M, Scherrer C, Stöckle R, Bachmann T (2020). "Designing Value Co-creation for a Free-Floating e-Bike-Sharing System." In H Yang, R Qiu, W Chen (eds.), *Smart Service Systems, Operations Management, and Analytics*, pp. 113–125. Springer International Publishing, Cham. ISBN 978-3-030-30967-1. doi:10.1007/978-3-030-30967-1_11.

Heitz C, Etschmann R, Stöckle R, Bachmann T, Templ M (2019). "User-based Redistribution in Free-floating Bike Sharing Systems." In *Operations Research Proceedings 2018*. Annual International Conference of the German Operations Research Society (GOR), Springer Nature Switzerland AG, Heidelberg, Germany. ISSN 0721-5924. `doi: 10.1007/978-3-030-18500-8{\_}69`.

Lippoldt K, Niels T, Bogenberger K (2019). "Analyzing the Potential of User-Based Relocations on a Free-Floating Carsharing System in Cologne." *Transportation Research Procedia*, **37**, 147 – 154. ISSN 2352-1465. `doi:https://doi.org/10.1016/j.trpro.2018.12.177`. 21st EURO Working Group on Transportation Meeting, EWGT 2018, 17th – 19th September 2018, Braunschweig, Germany.

Müller J, de Almeida Correia GH, Bogenberger K (2017). "An Explanatory Model Approach for the Spatial Distribution of Free-Floating Carsharing Bookings: A Case-Study of German Cities." *Sustainability*, **9**(7), 1–17. `doi:10.3390/su9071290`.

Reiss S, Bogenberger K (2016). "Validation of a Relocation Strategy for Munich's Bike Sharing System." *Transportation Research Procedia*, **19**, 341 – 349. ISSN 2352-1465. `doi: https://doi.org/10.1016/j.trpro.2016.12.093`. Transforming Urban Mobility. mobil.TUM 2016. International Scientific Conference on Mobility and Transport. Conference Proceedings.

Singla A, Santoni M, Bartók G, Mukerji P, Meenen M, Krause A (2015). "Incentivizing Users for Balancing Bike Sharing Systems." In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, p. 723–729. AAAI Press. ISBN 0262511290.

Templ M (2019). *mobi2: Demand Estimation of Vehicles in Free-sharing Systems*. R package version 0.1.1.

Wood S (2003). "Thin Plate Regression Splines." *Journal of the Royal Statistical Society Series B*, **65**(1), 95–114. `doi:10.1111/1467-9868.00374`.

Wood S (2017). *Generalized Additive Models*. Chapman & Hall/CRC. CRC Press, New York. `doi:10.1201/9781315370279`. Second edition.

Affiliation:

Matthias Templ, Christoph Heitz
Institute of Data Analysis and Process Design
Zurich University of Applied Sciences
CH-8400 Winterthur, Switzerland
E-mail: `matthias.templ@gmail.com`
URL: `https://www.zhaw.ch/de/ueber-uns/person/teml/`