

Article

Using Object Detection on Social Media Images for Urban Bicycle Infrastructure Planning: A Case Study of Dresden

Martin Knura ^{1,*}, Florian Kluger ^{2,†}, Moris Zahtila ^{3,†}, Jochen Schiewe ¹, Bodo Rosenhahn ² and Dirk Burghardt ³

- ¹ Lab for Geoinformatics and Geovisualization (g2lab), HafenCity University Hamburg, Henning-Voscherau-Platz 1, 20457 Hamburg, Germany; jochen.schiewe@hcu-hamburg.de
² Institute of Information Processing, Leibniz University Hannover, Appelstr. 9a, 30167 Hannover, Germany; kluger@tnt.uni-hannover.de (F.K.); rosenhahn@tnt.uni-hannover.de (B.R.)
³ Institute of Cartography, Dresden University of Technology, Helmholtzstr. 10, 01062 Dresden, Germany; moris.zahtila@tu-dresden.de (M.Z.); dirk.burghardt@tu-dresden.de (D.B.)
* Correspondence: martin.knura@hcu-hamburg.de
† These authors contributed equally to this work.

Abstract: With cities reinforcing greener ways of urban mobility, encouraging urban cycling helps to reduce the number of motorized vehicles on the streets. However, that also leads to a significant increase in the number of bicycles in urban areas, making the question of planning the cycling infrastructure an important topic. In this paper, we introduce a new method for analyzing the demand for bicycle parking facilities in urban areas based on object detection of social media images. We use a subset of the YFCC100m dataset, a collection of posts from the social media platform Flickr, and utilize a state-of-the-art object detection algorithm to detect and classify moving and parked bicycles in the city of Dresden, Germany. We were able to retrieve the vast majority of bicycles while generating few false positives and classify them as either moving or stationary. We then conducted a case study in which we compare areas with a high density of parked bicycles with the number of currently available parking spots in the same areas and identify potential locations where new bicycle parking facilities can be introduced. With the results of the case study, we show that our approach is a useful additional data source for urban bicycle infrastructure planning because it provides information that is otherwise hard to obtain.

Keywords: object detection; social media; urban planning; bicycle infrastructure; computer vision; volunteered geographical information; visual analytics



Citation: Knura, M.; Kluger, F.; Zahtila, M.; Schiewe, J.; Rosenhahn, B.; Burghardt, D. Using Object Detection on Social Media Images for Urban Bicycle Infrastructure Planning: A Case Study of Dresden. *ISPRS Int. J. Geo-Inf.* **2021**, *10*, 733. <https://doi.org/10.3390/ijgi10110733>

Academic Editor: Jean-Claude Thill, Ran Tao, Zhaoya Gong and Wolfgang Kainz

Received: 17 September 2021
Accepted: 23 October 2021
Published: 28 October 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Today, as cities grow and develop at high-speed rates, many of them actively reinforce greener ways of urban mobility to fight against pollution, traffic jams, noise, etc. [1]. One of the encouraged ways of urban commuting is cycling for its beneficial effect on both the environment and personal health [2]. While this strategy helps to reduce the number of motorized vehicles on streets, it also leads to a significant increase in the number of bicycles in urban areas, which raises the importance of bicycle infrastructure planning as a topic. Cities usually pay attention to providing a larger number of bicycle parking at popular locations such as train stations, shopping malls and highly frequented squares; however, numerous bicycles are often randomly parked in the surrounding areas where there are fewer bicycle racks available [3]. Accordingly, planning the cycling infrastructure is an important topic for both urban planners and cyclists.

Traditional methods of collecting field data in urban planning are spot observations and surveys. Spot observations are usually conducted by counting objects (e.g., passengers, bicycles, etc.) at urban locations of interest. This method is resource-consuming regarding time and staff, so the collected data usually do not cover longer time intervals. In contrast,

conducting surveys usually implies recruiting passengers and asking them to answer a questionnaire, which makes it difficult to collect a high number of answers and introduces a bias towards citizens who have a positive general attitude towards participation in a survey. In both cases, spot observations and surveys, there is a clear gap that concerns the availability of the data available to urban planners, and thus also for planning the bicycle infrastructure.

To tackle the topic of cycling in cities and improve the quality of information that urban planners use for making decisions related to spatial investments, researchers largely turned to newly available sources of data. The majority of this research focuses on analyzing the data from bicycle-sharing systems (BSS). The related data sources are mainly the pooling stations, which allow analyses based on the numbers of available bicycles and free parking spots at BSS stations [4] or check-ins and check-outs on BSS stations [5–7]. While being valuable for the logistics of the BSSs, these numbers do not necessarily give insights useful to planning the infrastructure for citizens who commute by using privately owned bicycles. Planning of parking facilities for bicycles that do not belong to the BSS is therefore not feasible with these data alone. Another inexpensive method of data collection is GPS tracking of ridden bicycles (e.g., via smartphone), which is, however, more suitable for analyzing trail patterns [8]. Some research related to urban planning also analyzed the mobile-phone data generated by mobile networks [9]. While collecting telecommunications activity can provide an extensive dataset, it does not differentiate bicycle users from other passengers.

In this paper, we address the gap in data collection for urban planning by focusing on social media data. Because of the continuous increase in the numbers of smartphone owners and social media users, we identify another opportunity to collect the needed bicycle-related data and develop a novel method for analyzing the demand for bicycle parking infrastructure in urban areas. We propose to use this method alongside others established in urban planning in order to enrich the data coverage and provide more comprehensive information for making decisions related to urban infrastructure investments, e.g., bicycle parking. We start in Section 2 with the hypothesis that social media posts can be useful for analyzing locations in cities in relation to bicycle usage. Section 3 introduces our method and data used to detect bicycles on photos from social media posts, and Section 4 presents the preliminary results of the bicycle detection process. Within our case study in Section 5, we show that our data processing can provide substantial value for planning bicycle parking facilities in the city of Dresden and discuss advantages and drawbacks of our approach in Section 6 before concluding in Section 7.

2. Related Work

For the most effective promotion of cycling in a city, planning bicycle infrastructure should be demand-driven, and so urban planners need to know the main characteristics of the bicycle traffic flows in their city [10]. Although there is an increasing interest in bicycles as part of multi-modal urban infrastructure, bicycle-related research in recent years focused mainly on aspects such as bicycle safety [11], positive impacts of cycling on public health [12], travel mode choice [13] and route choice analysis [10,14]. By contrast, less attention is paid to traffic engineering topics such as traffic counts, travel times, and capacities [15].

In general, bicycle traffic volume data are hard to obtain. As opposed to motorized traffic, bicycle traffic volume is strongly affected by the presence and qualities of bicycle infrastructure, elevation, motorized vehicles, weather conditions, etc. [16]. Historically, research on cycling activity relied on individual-level surveys on household travel—methods that are resource intensive and can produce statistically unrepresentative samples distorting the findings of the qualitative analysis [17]. Modern methods of bicycle traffic estimation fall into two categories: long-term counters that run continuously, and short-term measurements of typically 1 to 28 days. To derive robust demands based on shorter observation periods, the values can be multiplied by scaling factors and factor groups

accounting for daily, weekly and seasonal bicycle volume variance gained through the continuous measurements [18].

In practice, the actual counting and tracing of bicycles can be executed using different data collection methods, which should follow quality assurance procedures [19]. These methods include:

- Use of stationary sensors to count passing bicycles,
- Analysis of public surveillance videos through object detection,
- GPS-tracking through devices used by cyclists,
- Tracking of GPS devices directly mounted on bicycles.

Adapting traditional methods for motor vehicle traffic monitoring, numerous technical solutions and commercial products to count cyclists with stationary sensors are available. For example, ref. [20] used data from pneumatic tubes on streets and radio beams on cycle paths for their study, while the data of [18] were obtained from inductive loop counters. Similar to sensors, the visual detection of bicycles using stationary cameras is adapting well-established techniques, in this case from computer vision [15,21], and can provide vehicle detection, classification, counting and speed measurements in real-time [22].

In contrast to stationary sensors or cameras, GPS tracking devices are capable of collecting data from a complete journey and can be divided into two categories, depending on the method of obtaining the data. First, the GPS data can be obtained through a device used by the cyclists, e.g., a smartphone application. These data are normally shared from volunteers either for scientific research [10,23], or for commercial use through mobile apps such as Strava [11]. Second, the tracking device can be directly mounted onto the bicycle. Dockless BSS provides real-time GPS data for every bicycle, which provides detailed insights into bicycle-sharing users' temporal and spatial mobility patterns [24,25]. For example, ref. [16] used GPS data of BSS provided by the company Wavelo.

Stationary sensors, traffic surveillance cameras and GPS tracking devices differ not only in the method of obtaining information but also in the characteristics of the data they provide, namely the gained spatial information, the coverage of target groups and the surveyed bicycle state. In general, stationary sensors have the advantage of covering every single passing bicycle, while they are fixed to a defined location and therefore only provide point-related data of moving bicycles. Visual analysis of traffic videos can identify *moving* and *parked* bicycles within the frame and is able to generate trajectory data within the covered area when cyclists are tracked over consecutive frames. While both types of GPS tracking devices provide trajectory data of the whole trip, a major disadvantage of this data collection method is the creation of biased "voluntary response samples", because it only includes data of people who have chosen to volunteer [26]. Furthermore, as the tracking device is not mounted onto the bicycle, the status and position of the bicycle while the bicyclist is not using the bicycle is unknown. In contrast, the status and position of integrated GPS devices can constantly be measured, which allows detecting the location where the bicycle is parked. However, unlike in Asia, urban mobility planning policies in Europe focus on *private* bicycle use [27], and *public* bicycles from BSS are more frequently used for first- and last-mile connection and leisure activities and less frequently for commuting [24,28].

As shown above, obtaining data on parked bicycles is still challenging. Information retrieval using social media data can be a complementary way of data collection. Social media usage is widespread geographically as well as temporally and has become a natural part of people's daily lives. As a consequence, data are generated implicitly by the users, providing an "in-the-wild sensing" of the city without restrictions of laboratory environments [29]. Social media posts usually contain text and time information, with potentially more visual (images, videos) and spatial data attached, which allows location extraction [30,31].

While there are several approaches to recognize low-level (e.g., walking, sitting, etc.) and high-level (e.g., eating, shopping, etc.) activities mainly based on different sources of social media data [32,33], we want to extract bicycle-related information solely using

images from social media posts. Regarding the identified data characteristics of bicycle-related measures stated above, this has two advantages. First, using images from social media potentially allows us to cover the whole area of the city, depending on the frequency of posts, and obtain information on a larger variety of bicycle usage. Second, we can distinguish between moving and parked bicycles using similar object detection methods as implemented for stationary traffic surveillance videos.

3. Method

Our approach for counting bicycles in images from social media posts consists of two steps. First, we applied a state-of-the-art object detection algorithm (Section 3.2) in order to detect and localize *bicycles* and *persons* in each image. Using the detected persons, we then classified each detected bicycle as either *moving* or *stationary* (Section 3.2.1). For evaluation (Section 3.3) and parameter selection (Section 3.4), we furthermore labeled an appropriate dataset (Section 3.1).

3.1. Dataset

In order to quantitatively evaluate the feasibility of using social media data for bicycle traffic analysis, we used the YFCC100m [34] dataset because it is one of the largest open-source datasets of its kind with a collection of 100 million posts from the social media site Flickr. Each post contains an image or video as well as additional information, such as location, time of capture and tags. All images were taken in the years between 2004 and 2014 and are scattered across the whole world. As we are mainly interested in data from urban areas, we selected a subset of images taken in a single city. This subset contains 30,922 images with location metadata indicating that they were recorded in the city of Dresden, Germany.

3.1.1. Bicycle Annotations

We manually annotated all bicycles in the subset of images. Each bicycle is labeled with a bounding box and assigned one of two categories: *stationary* if the bicycle is currently parked, or *moving* if it is being ridden, wheeled or otherwise in use. Of the 30,922 images, 2219 (7.2%) contain at least one bicycle, with 1457 (4.7%) images containing stationary bicycles and 976 (3.2%) images containing moving bicycles. As Figure 1 shows, most images (1204, 54.3%) contain exactly one bicycle. However, images with significantly larger numbers of bicycles occur as well, e.g., 100 images (4.5%) depict more than six bicycles. In total, we labeled 4913 bicycles, of which 3038 (61.8%) are stationary and 1875 (38.2%) are moving. Figure 2 shows a few examples.




Figure 1. This histogram shows the numbers of images in the Dresden subset of the YFCC100m dataset containing between one and six, and more than six bicycles.




Figure 2. Examples from our annotated Dresden subset of the YFCC100m dataset. We manually labeled both *moving* (yellow boxes) and *stationary* (cyan boxes) bicycles.

3.2. Object Detection

In order to automatically and reliably count the number of moving and stationary bicycles in an image, we utilized a state-of-the-art object detection algorithm. The task of object detection comprises localization of objects in the image, usually by estimating the coordinates of bounding boxes framing the objects, as well as classifying each object using a set of predefined categories. Numerous approaches for object detection have been presented in recent years [35–39]. They all use convolutional neural networks (CNNs) and are trained on the large-scale COCO (Common Objects in Context) [40] dataset. COCO contains more than 200,000 images labeled with object bounding boxes of 80 different categories such as car, bicycle, person, couch, orange, etc. For all experiments in this work, we used the recently presented EfficientDet [35] object detection algorithm, which has been pre-trained on the COCO dataset, as it provides state-of-the-art performance. Compared to the previous best method [41], EfficientDet achieves a significantly higher mean average precision (mAP) on the challenging COCO dataset (54.4% vs. 50.7%) while being computationally more efficient. Computing object detections for one image on an Nvidia Titan V GPU takes 285 ms with EfficientDet, while [41] requires 489 ms, i.e., almost twice as long.

Given an image I , the object detection algorithm computes a set \mathcal{P} of object proposals $\mathbf{P}_i = (\mathbf{b}_i, c_i, s_i) \in \mathcal{P}$. Each object proposal is defined by a bounding box (rectangle) \mathbf{b}_i with image coordinates $[x_{i,1}, y_{i,1}, x_{i,2}, y_{i,2}]$, an object class c (e.g., *bicycle*), and a confidence score, s which can be loosely interpreted as an estimate of the likelihood that the object proposal is correct. In practice, object proposals that have a confidence score below a threshold θ_s are discarded. This threshold must be chosen appropriately in order to minimize the number of false detections while maximizing the number of correct detections.

In the following, we are only interested in bicycle detections $\mathbf{P}_b \in \mathcal{P}_b \subseteq \mathcal{P}$ and person detections $\mathbf{P}_p \in \mathcal{P}_p \subseteq \mathcal{P}$, with $\mathcal{P}_b \cap \mathcal{P}_p = \emptyset$:

$$\begin{aligned} \forall \mathbf{P}_i \in \mathcal{P} : (c_i = \text{bicycle} \Leftrightarrow \mathbf{P}_i \in \mathcal{P}_b) \\ \wedge (c_i = \text{person} \Leftrightarrow \mathbf{P}_i \in \mathcal{P}_p). \end{aligned} \quad (1)$$

3.2.1. Moving Bicycles

In order to differentiate between moving and stationary bicycles, we leverage the ability of the object detector to localize people in addition to bicycles. We assume that

if a bicycle is located right below a person or right next to a person, this bicycle is being handled by that person and is thus non-stationary or *moving*. In that case, the center of the bounding box of a detected person \mathbf{P}_p must be located above the bounding box center of bicycle \mathbf{P}_b . We describe this relation via the following indicator function:

$$\chi(\mathbf{P}_b, \mathbf{P}_p) = \begin{cases} 1 & \text{if } y_{p,1} + y_{p,2} > y_{b,1} + y_{b,2}, \\ 0 & \text{else.} \end{cases} \quad (2)$$

Since a person must be located in very close proximity to a moving bicycle, we assume a minimal overlap of their respective bounding boxes. We measured this overlap using the intersection-over-union (IoU) metric, which computes the ratio of the overlapping area of the bounding boxes to their unified area:

$$\text{IoU}(\mathbf{P}_b, \mathbf{P}_p) = \frac{A(\mathbf{b}_b \cap \mathbf{b}_p)}{A(\mathbf{b}_b \cup \mathbf{b}_p)} \in [0, 1]. \quad (3)$$

For every bicycle detection $\mathbf{P}_b \in \mathcal{P}_b$ and every person detection $\mathbf{P}_p \in \mathcal{P}_p$, we define an overlap matrix \mathbf{C} with:

$$C_{bp} = \chi(\mathbf{P}_b, \mathbf{P}_p) \cdot \text{IoU}(\mathbf{P}_b, \mathbf{P}_p). \quad (4)$$

Using the Hungarian method [42], we find a maximum overlap assignment \mathcal{H} based on \mathbf{C} . If a bicycle \mathbf{P}_b is assigned to a person \mathbf{P}_p with $C_{bp} > \theta_p$, we define the bicycle as *moving* and as *stationary* otherwise:

$$\begin{aligned} \forall \mathbf{P}_i \in \mathcal{P}_b : ((\exists \mathbf{P}_j \in \mathcal{P}_p ([i, j] \in \mathcal{H} \wedge C_{ij} > \theta_p)) \Leftrightarrow \mathbf{P}_i \in \mathcal{P}_{bm}) \\ \wedge (\mathbf{P}_i \notin \mathcal{P}_{bm} \Leftrightarrow \mathbf{P}_i \in \mathcal{P}_{bs}), \end{aligned} \quad (5)$$

with \mathcal{P}_{bm} and \mathcal{P}_{bs} denoting the sets of moving and stationary bicycle detections, respectively. Figure 3 shows a few examples of bicycles that have been classified as *stationary* or *moving* using this procedure. We denote the maximum assigned overlap with a person for a bicycle detection \mathbf{P}_b as $C_b = \max_p C_{bp}$.

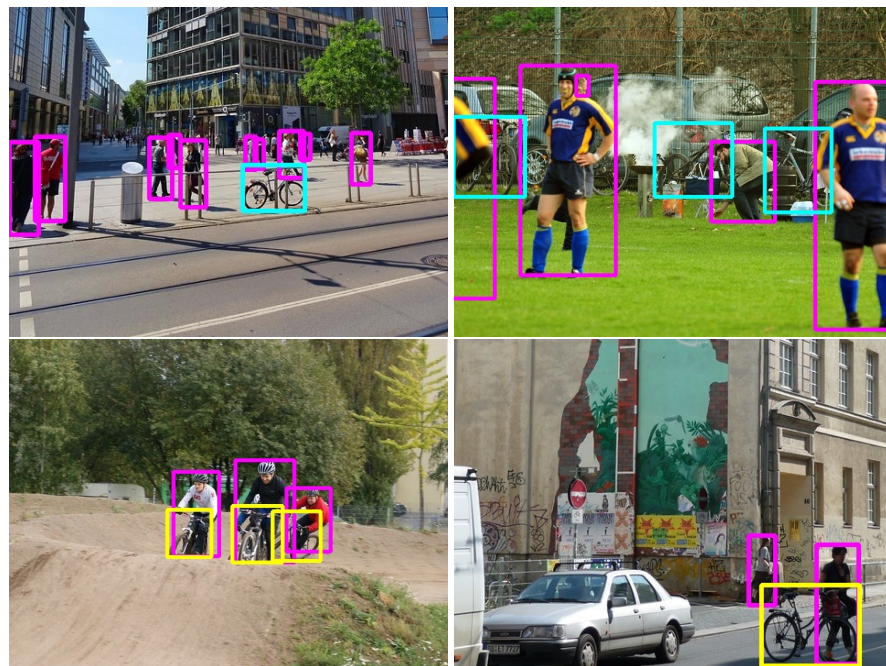


Figure 3. Examples of correctly identified stationary (top, cyan boxes) and moving (bottom, yellow boxes) bicycles, with all detected persons marked in magenta boxes.

3.3. Evaluating Detections

In order to evaluate the bicycle detection method and to optimize its parameters, we compared the proposed bicycle detections with the ground truth annotations (cf. Section 3.1.1). Given ground truth annotations $\mathbf{T}_i = (\hat{\mathbf{b}}_i, \hat{c}_i, \hat{s}_i) \in \mathcal{T}$ and proposed detections $\mathbf{P}_i \in \mathcal{P}$ for the same image, we define an overlap matrix \mathbf{D} with:

$$D_{ij} = \text{IoU}(\mathbf{P}_i, \mathbf{T}_j) \forall \mathbf{P}_i \in \mathcal{P}, \mathbf{T}_j \in \mathcal{T}. \quad (6)$$

We find a maximum overlap assignment based on \mathbf{D} using the Hungarian method [42]. If a prediction \mathbf{P}_i is assigned to an annotation \mathbf{T}_j with $D_{ij} > \theta_{\text{IoU}}$ and same object class $c_i = \hat{c}_j$, we regard it as a true positive $\mathbf{P}_i \in \mathcal{P}_{\text{tp}}$. Otherwise, it is a false positive $\mathbf{P}_i \in \mathcal{P}_{\text{fp}}$. Likewise, if an annotation \mathbf{T}_j is not assigned to a prediction, it counts as a false negative $\mathbf{T}_j \in \mathcal{T}_{\text{fn}}$. As localization accuracy is of little relevance for our application—we only need to know the number of bicycles in an image—we set the IoU threshold relatively low, i.e., $\theta_{\text{IoU}} = 0.1$. After assigning predictions and annotations for each image, we can compute precision and recall over all images in order to assess the quality of the predictions. *Precision* is defined as the ratio of the number of correctly detected objects (true positives) to the number of all detections (true positives and false positives):

$$\text{precision} = \frac{|\mathcal{P}_{\text{tp}}|}{|\mathcal{P}|} = \frac{|\mathcal{P}_{\text{tp}}|}{|\mathcal{P}_{\text{tp}}| + |\mathcal{P}_{\text{fp}}|}. \quad (7)$$

Recall is the ratio of the number of correctly detected objects (true positives) to the number of all present objects (true positives and false negatives), i.e., all annotated objects:

$$\text{recall} = \frac{|\mathcal{P}_{\text{tp}}|}{|\mathcal{T}|} = \frac{|\mathcal{P}_{\text{tp}}|}{|\mathcal{P}_{\text{tp}}| + |\mathcal{T}_{\text{fn}}|}. \quad (8)$$

We purposely do not use the mean average precision metric (mAP, cf. Section 3.2) commonly utilized in object detection literature for evaluation with the COCO dataset [35,41]. The mAP metric computes the mean of the area under the precision-recall curve over a range of $\theta_{\text{IoU}} \in [0.5, 0.95]$. While this metric is well suited for comparing the performance of object detection algorithms independent of confidence threshold θ_s and partially independent of IoU threshold θ_{IoU} , it does not provide information about the accuracy of an algorithm in a practical setting, where these thresholds must be set to a specific value.

3.4. Determining Thresholds

We empirically determine a confidence threshold θ_s and a person assignment threshold θ_p in order to strike an optimal balance between precision and recall.

3.4.1. Confidence Threshold

We adjust precision and recall for all bicycle detections—both moving and stationary—by changing the confidence threshold θ_s . We compute recall and precision for all values of $\theta_s \in [0, 1]$ and show the results in Figure 4. The first graph in Figure 4 shows corresponding recall and precision values, and the second and third graphs show recall and precision values corresponding to different threshold values. We identify a point on the recall-precision curve which is as close to the top-right corner as possible, i.e., maximizing both precision and recall. This point corresponds to a threshold of roughly $\theta_s = 0.4$, resulting in a precision of 0.96 and recall of 0.81.




Figure 4. Precision and recall of bicycle detections in relation to the confidence threshold θ_s : the first graph shows the recall–precision curve, while the second and third graphs represent the relationships of recall and precision to the confidence threshold separately.

3.4.2. Person Assignment Threshold

In order to determine an optimal person assignment threshold θ_p , we considered all true positive bicycle detections $\mathbf{P}_b \in \mathcal{P}_{tp}$ and their assigned maximum person overlap C_b . For all thresholds $\theta_p \in [0, 1]$, we computed the fraction of detections that are correctly classified as either *moving* or *stationary*. As Figure 5 shows, this classification accuracy peaks at roughly 89.5%. We thus set the person assignment threshold to the corresponding value of $\theta_p = 0.15$.




Figure 5. Moving vs. stationary: we plot the classification accuracy for a range of person assignment thresholds θ_p in order to identify an optimal value.

4. Bicycle Detection Results

4.1. Detection Accuracy

In order to assess the overall accuracy of our approach, we compare our bicycle detections with the ground truth bicycle annotations from our dataset (cf. Section 3.1). As the confusion matrix in Table 1 shows, we detected a total of 4157 bicycles in Dresden, from which 1589 were classified as moving and 2568 as stationary. Of these 4157 detections, only 160 were incorrect, resulting in a false discovery rate of 3.85% and equivalently a precision of 96.1%. We correctly identified 3997 of the 4913 bicycles in the dataset, thus achieving a recall of 81.4%. This means that we have adjusted our bicycle detection method to operate rather cautiously, i.e., the number of false positives is significantly lower than the number of false negatives. The vast majority of false negatives can be divided into three categories: small (i.e., low resolution) bicycles, partly occluded bicycles, and unusual perspectives. Figure 6 shows one example image for each category. In such cases, the bicycles may be difficult to recognize even for a human annotator.

Table 1. This confusion matrix shows the number of bicycles of certain ground truth classes (rows) being classified into estimated classes (columns) by our method. The Σ -entries indicate column- and row-wise sums. *None* indicates either no bicycle present or no corresponding bicycle detected and were omitted from the overall sums.

True \ Estimated	Moving	Stationary	None	Σ
moving	1368	226	281	1875
stationary	194	2209	635	3038
none	27	133	-	(160)
Σ	1589	2568	(916)	4913
				4157



Figure 6. Most common cases of false negatives, i.e., unidentified bicycles, from left to right: small size or unfavorable lighting conditions, partial occlusions, unusual pose of bicycle or camera.

The smaller number of false positives fall into the following four categories: parts of complete bicycles (i.e., possibly duplicates), other wheeled objects (such as motorcycles, wheelchairs or baby strollers), traffic signs with bicycle pictograms, and miscellaneous. We present one example of each kind in Figure 7.




Figure 7. Most common cases of false positives, i.e., wrongly detected bicycles, from left to right: smaller parts of complete bicycles; other wheeled objects such as baby strollers, wheelchairs and motorcycles; traffic signs with bicycle pictograms; miscellaneous objects such as musical instruments, chairs and camera tripods.

Within the set of correctly identified bicycles, we classify most of the moving bicycles (1368 of 1594, 85.8%) and most of the stationary bicycles (2209 of 2403, 91.9%) correctly. False classifications as *moving* most commonly occur when a person is coincidentally located in close proximity to a stationary bicycle, or when such a person is falsely detected. False classifications as *stationary* occur when the overlap between the bicycle and person detections is too small, when the person was not detected at all, or when the bicycle is moving without a person (e.g., mounted on a car). We provide one example for each case in Figure 8.




Figure 8. Examples of false classifications into the *moving* (yellow) and *stationary* (cyan) categories. Person detections are marked in magenta. From left to right: person coincidentally located next to a stationary bicycle; false person detection; small overlap between person and bicycle detections; undetected person; bicycle moving without a person.

Bicycles were detected in 2058 different images, with more than half of these images containing only one bicycle (Figure 9). In general, the number of bicycles per image follows a similar distribution as shown in Figure 1 but with slightly fewer images with more than one bicycle. In the following subsection, we describe the spatial distribution of the detected bicycles in Dresden and compare our results with other bicycle-related data sources.




Figure 9. All locations of images in Dresden containing at least one bicycle detection, and a histogram of the number of bicycle detections per image.

4.2. Spatial Distribution of Recognized Bicycles

Figure 9 shows the spatial distribution of the images containing at least one bicycle in the area of Dresden. There are a few major patterns visible within this dataset. The majority of the images is located in the inner city of Dresden, starting from the central train station in the southwest and following the main shopping mile towards the old town (*Altstadt*) and the Elbe river. Two smaller clusters of bicycles surround this main cluster: one on the northern side of the river within the new town (*Neustadt*), and another one at the university

campus south of the main station. Furthermore, a significant number of images are located along the Elbe river, on a path that is also a part of the Elbe Cycle Route that follows the river from the border between Germany and Czechia to the river's mouth at the North Sea. Smaller clusters are related also to traffic nodes such as bridges and public transport hotspots, and to the city's largest park, Großer Garten.

Of the 2058 images with bicycles detected, 1300 contain at least one stationary bicycle, and 1044 contain at least one moving bicycle. Therefore, 286 images contain both stationary and moving bicycles. In the following, we present the distributions of each type over the area of Dresden and elaborate on their differences.

Moving bicycles: Only one moving bicycle was detected in 923 of the 1044 images (88.4%). While the overall pattern follows the same distribution as the whole dataset described above, images with multiple moving bicycles mostly occur in the Altstadt and along the Elbe river (Figure 10a). All of the six images with more than eight moving bicycles are located on the two small bridges crossing the Elbe.

Stationary bicycles: Similarly to moving bicycles, the vast majority of the 1300 images with stationary bicycles contain only one bicycle (79.8%), and the overall pattern is the same as for the whole dataset (Figure 10b). In contrast to moving bicycles, there are significantly more stationary bicycles in the northern part of Neustadt (*Äußere Neustadt*) and on the university campus. Most of the images with more than eight bicycles are located either at train stations or on the campus.

Mixed images: The distributions of stationary and moving bicycles show some small differences in the locations of multiple bicycle detections, so we decide to have a closer look at the 286 images that contain both moving and stationary bicycles. Figure 10c shows the locations of these images and the respective majority class. Again, similar patterns as described above are visible: images with more stationary bicycles are located at the train stations, and images with more moving bicycles are located at the cycling path along the Elbe and on the bridges.




Figure 10. Spatial distributions of classified bicycle types. (a) Moving bicycles; (b) stationary bicycles; (c) images with both moving and stationary bicycles. Scale and orientation are equal for all maps.

4.3. Density of Images with Detected Bicycles

As volunteered geographic data show an enormous spatial heterogeneity, not only the total number of detected bicycles is of major interest but also the portion of images containing bicycles in relation to the whole dataset in the same area. We argue that our analysis for a certain area is more meaningful if more images are available for that area in our dataset. Furthermore, the more frequently bicycles appear in images within a certain area of the city, the more important that area is with regards to the focus of our paper. Therefore, we create grids with various cell sizes—each square cell covering a small area of the city—and sum up for each cell area:

- The number of images in the YFCC100m Dresden subset,
- The number of bicycles detected on images from the subset,
- The number of stationary bicycles detected,

- The number of moving bicycles detected,
- The number of images containing at least one bicycle,
- The number of images containing at least one stationary bicycle,
- The number of images containing at least one moving bicycle.

With these data, we can calculate the densities of recognized bicycles for every grid cell. Figure 11 shows the result of processing with grid cells of 100×100 m, as we combine the total number of images with the percentage of images containing a bicycle in a bivariate map. We can identify the same pattern as in Figure 9 but now with the additional information of the total number of images. Based on this, we conclude that our dataset has more significance for the inner city of Dresden, as well as for a few particular spots along the Elbe river and on the university campus. With this visualization, we can also identify a true outlier southeast of the inner city, combining a high number of images with a high percentage of bicycle-containing images. As we found out, this grid cell contains a bicycle race track facility (Bike Areal Dresden), which confirms that our method is able to detect different levels of bicycle activity in the city.




Figure 11. The map shows a combination of the total number of images with the percentage of images containing a bicycle for each grid cell with a size of 100×100 m.

4.4. Comparison with Other Datasets

After introducing our method and the distribution of bicycle-containing images in Dresden, we now compare our dataset to other relevant sources for bicycle-related traffic (Figure 12):

- The positions of the stationary bicycle counting sensors installed by the city administration of Dresden.
- The defined bicycle return areas and a sample of positions of parked bicycles from the bicycle sharing system MOBI, obtained at different points in time throughout a week in September 2021.




Figure 12. Positions of bicycle counting stations and the different return zones of MOBI bicycles in Dresden, where returning a bicycle at MOBIpoints and on blue streets is free of charge, while returning on a pink zone is charged with 1 Euro and outside the other areas with 20 Euro.

4.4.1. Bicycle Counting Stations

Until 2021, the city administration of Dresden built nine stationary sensors at eight different places in the city using pneumatic tubes, as mentioned in Section 2. Figure 12 shows the locations of the counters, which were distributed over the city based on the different criteria such as centrality, route characteristics (cycle route, mixed traffic, road crossings) and the assumed type of destination or purpose of the route (work, shopping, leisure) [43]. Compared to the characteristics of our dataset, these counters cover different areas, as they are all located in areas where our dataset only provides a few images and detected bicycles. Five of the stations are placed in a 100×100 m grid cell where we do not find a single image with a bicycle. If we use larger grid cells of 400×400 m, all of the stations are located in grid cells containing at least three detected bicycles. We argue that these differences originate from the differing purposes of data collection: the municipal bicycle counters are planned to measure the flows of cyclists, particularly between the suburbs and the city center, whereas our method provides the positions of bicycles as an *in-the-wild sensing*, mainly in the city center.

4.4.2. Bicycle Sharing Systems

A joint venture of the Dresden public transport services (DVB) and the company Nextbike offers a dockless bicycle sharing system named MOBI in the city of Dresden. There are different bicycle return zones for their bicycles defined in the city: blue streets, which cover the main roads in the inner city of Dresden as well as in the inner suburbs; pink zones, which cover the areas in the inner city outside the blue streets; and the area

outside the zones, which covers all other areas in Dresden (Figure 12). The return fee users have to pay after using a bicycle depends on the type of zone they return it in. Returning a bicycle on a blue street is free of charge, while returning a bicycle in a pink zone or outside yields service fees of 1 and 20 Euros, respectively. In addition, parking stations—so-called MOBIpoints—are defined at certain locations (also in the outside of the zones), where it is possible to return a bicycle free of charge and earn 10 free minutes for the next rent. The real-time positions of all parked bicycles can be requested via an API (nextbike API for real-time locations. https://api.nextbike.net/api/documentation#maps_api_-_locations, accessed on 8 September 2021), which makes it possible to obtain a dataset of all parked MOBI bicycles within the city. We retrieved positions for five different days in September 2021, and collected a total number of 3160 locations of parked MOBI bicycles in Dresden: 79.8% of the bicycles were parked on blue streets and at MOBIpoints, 10.4% in the pink zones, and 9.7% outside of the zones. In general, the distribution of MOBI bicycles covers the same area as our dataset, but due to the pricing system, the majority of bicycles are returned on the blue streets. Thereafter, the dataset can be described as a *biased in-the-wild sensing* where the flows represent the bicycle sharing demand, but the actual positions of parked bicycles result from the pricing incentives.

5. Case Study

With the Flickr dataset, we focus on analyzing the general situation in Dresden related to parked bicycles. We consider the bicycles classified as stationary as parked bicycles and use index PB to refer to them. To take into account the credibility of the data in urban areas, we perform several steps of analysis based on a visual comparison of data. We start by comparing the number of bicycle detections that were classified as stationary (N_{PB}) with the percentage of photos that contain detections of stationary bicycles in relation to all photos posted in that area (P_{PB}). Next, we compare the number of detected parked bicycles N_{PB} with the number of currently available parking spots (N_{PS}). In the analysis, we use 100×100 m grid cells that represent units of the urban area, above which we visually overlap layers of N_{PB} , P_{PB} , and N_{PS} . Finally, we detect the most critical areas in the city of Dresden in regards to the planning of bicycle parking facilities.

5.1. Number of Parked Bicycles N_{PB} vs. Percentage of Photos Containing Parked Bicycles P_{PB}

We visually analyze each combination of layers' classes (as shown in Figure 13). The classes were created by slightly adapting the natural breaks (Jenks [44]) classification. For N_{PB} , we distinguish five classes of the occurring number of detections: 1 to 3—low number of bicycle detections, does not require immediate attention; 4 to 8—medium number of bicycle detections, requires monitoring; 9 to 18—high number of bicycle detections, requires analysis; 19 to 35—high number of bicycle detections, urgently requires analysis; and 36 to 65—very high number of bicycle detections, requires immediate analysis. For P_{PB} we classify the occurring detection percentages into four classes: 0.1 to 3.0%—very low, insignificant; 3.1 to 8.0%—low, moderately significant; 8.1 to 18.0%—medium, significant; and 18.1 to 35.0%—medium-high, highly significant.

The existence of *grid overlaps* of N_{PB} and P_{PB} shows us the occurrence of different class combinations. To better understand the importance of those combinations for analysis, we assign each combination an expected relevance in planning the bicycle parking (*low*, *medium*, or *high relevance*) depending on previously defined class importance (Table 2). For example, we expect the grid cells where a higher percentage of photos contains a higher number of bicycle detections to be significantly more relevant for planning the bicycle parking than grid cells with a low percentage of photos and a low number of bicycle detections. The designation of relevance will help us in later steps to prioritize cells of urban space to the need for further analysis using additional datasets.

Table 2. Relevance assigned to different combinations of the percentages of photos in the dataset containing stationary bicycles P_{PB} and the number of stationary bicycle detections on photos N_{PB} : low (grey 20%), medium (grey 55%), and high (grey 85%) relevance.

N_{PB} \ P_{PB} (%)	0.1–3.0	3.1–8.0	8.1–18.0	18.1–35.0
1–3	Low	Low	Low	Low
4–8	Low	Medium	Medium	Medium
9–18	Low	Medium	High	High
19–35	Medium	High	High	High
36–65	Medium	High	High	High

Next, for each N_{PB} and P_{PB} class combination, we counted the number of cells where the two layers overlap, i.e., when a cell falls into a certain combination of N_{PB} and P_{PB} categories. The results (Table 3) show that some combinations have no overlaps, while some have an overly high number of overlaps. For example, highly relevant cells with P_{PB} of 18.01 to 35.0% and 9 to 18 detections have zero occurrences, while cells of low relevance with P_{PB} of 0.01 to 3.0% and 1 to 3 detections provide a number of overlaps so high that it is not meaningful to precisely count and analyze them visually.

Lastly, for existing overlaps, we analyze in which urban areas they mainly occur (Table 3). Most of the highly relevant combinations occur in the city center of Altstadt, while cells of medium and low relevance are distributed all over the city, including Altstadt. In the following step, we analyze grid cells and detect critical locations in the city that require further attention. We achieve this by introducing the number of currently available parking spots N_{PS} .

Table 3. Number and common urban location of 100×100 m grid cells in Dresden for each combination of the photos in the dataset containing stationary bicycles P_{PB} and the number of stationary bicycle detections on photos N_{PB} . Shades of gray refer to relevances assigned in Table 2.

N_{PB} \ P_{PB} [%]	0.1–3.0	3.1–8.0	8.1–18.0	18.1–35.0
1–3	No. of cells: >100 Location: All over the city (both city center and outside)	0	0	0
4–8	61 Mainly in the city centers of Altstadt and Neustadt, university campus and Großergarten	31 Mainly in the city centers of Altstadt and Neustadt	0	0
9–18	8 Outside of the city center of Altstadt	29 Mainly in the city center of Altstadt	7 Mainly in the city center of Altstadt	0
19–35	0	0	17 Mainly in the city center of Altstadt and main train station	1 City center of Altstadt
36–65	0	1 Main train station	5 Mainly at big train stations and university library	3 City center of Altstadt

5.2. Number of Parked Bicycles N_{PB} vs. Percentage of Photos Containing Parked Bicycles P_{PB} vs. Number of Available Parking Spots N_{PS}

After identifying for which combinations of N_{PB} and P_{PB} overlaps exist and where, we further analyze the overlaps in relation to the number of bicycle parking spots N_{PS} (data downloaded from Open Street Map via the Overpass API. <https://overpass-turbo.eu/>, accessed on 5 June 2021) in order to detect locations of a *potential parking space deficit*. For this purpose, we first classify grid cells into three categories: (I) Cells contain sufficient parking capacity; (II) Cells partially contain parking capacity; (III) Cells contain no parking capacity; and one subcategory: (a) Cells have close access to the neighboring cell's parking capacity.

We classify each cell into one of the three categories of parking capacity based on the relation of N_{PB} and N_{PS} for that cell (Table 4). As N_{PB} and N_{PS} are quantized into sets of value ranges, we consider the upper bounds for each parking facility. For example, if N_{PS} is in the range of 12 to 37, we assume $N_{PS} = 37$. If the upper bound of N_{PS} equals or exceeds the upper bound of N_{PB} in the analyzed cell, then we assume that the cell contains sufficient parking capacity (I). If the upper bound of N_{PS} is lower than the upper bound of N_{PB} for a cell, we consider that the cell partially includes parking capacity (II). Otherwise, we assume that the cell contains no parking capacity (III). For example, there are three cells with P_{PB} of 18.01 to 35.0% containing 36 to 65 parked bicycles N_{PB} . Two out of three cells do not contain any parking ($N_{PS} = 0$), and one contains two parking facilities: one of capacity up to 12 bicycles and the other of capacity 12 to 37 bicycles (Figure 13). Since the upper bound of N_{PS} is 49, it does not reach the upper bound of N_{PB} of 65; in this manner, we consider the cell is partially provided with parking capacity.



Figure 13. Example of visual analysis of bicycle parking capacity in Dresden by overlapping 100×100 m grid cells of the N_{PB} (gray) and P_{PB} (hashed) layers in QGIS (version 3.16).

Considering that we implemented the above categorization to identify how *critical* the condition within the cell regarding the lack of parking spots generally is, we also introduced the subcategory *a* for the cells that have close access to the neighboring cell parking capacity. The motivation for this is the fact that we were able to identify cases where the cell belonged to a category II or III but parking of the neighboring cell was located at the exact border between its native cell and the analyzed cell. Therefore, the subcategory serves us to decide whether the cell we analyze is less critical because the parking can be easily reached outside the cell. For example, there are five cells with a P_{PB} of 8.1 to 18.0% containing 36 to 65 detections of parked bicycles N_{PB} . Two out of five cells contain a sufficient number of parking spots and three are partially covered with

parking capacity. However, two out of the three cells that are partially covered with parking capacity allow easy access to the neighboring cell's parking. We conclude that, even though the whole category is of high relevance, it is fairly well covered with parking facilities and, consequently, does not require immediate attention. The results of the analysis are presented in Tables 4 and 5.

Table 4. Number of 100 × 100 m grid cells that: contain sufficient parking capacity (I); are partially covered with parking capacity (II); are partially covered with parking capacity but some cells have close access to the neighboring cell's parking capacity (IIa); contain no parking capacity (III); or contain no parking capacity but some cells have close access to the neighboring cell's parking capacity (IIIa). Shades of gray additionally show relevances assigned in Table 2.

N_{PB} \ P_{PB} [%]	0.1–3.0	3.1–8.0	8.1–18.0	18.1–35.0
1–3	/	/	/	/
4–8	I: 30/61 II: 0/61 III: 31/61; IIIa: 6/31	I: 14/31 II: 0/31 III: 17/31; IIIa: 1/17	/	/
9–18	I: 3/8 II: 1/8 III: 4/8; IIIa: 2/4	I: 15/29 II: 2/29 III: 12/29; IIIa: 1/12	I: 1/7 II: 2/7 III: 4/7; IIIa: 2/4	/
19–35	/	/	I: 6/17 II: 4/17; IIa: 1/4 III: 7/17; IIIa: 3/7	I: 0/1 II: 0/1 III: 1/1; IIIa: 1/1
36–65	/	I: 0/1 II: 0/1 III: 1/1; IIIa: 1/1	I: 2/5 II: 3/5; IIa: 2/3 III: 0/5	I: 0/3 II: 1/3 III: 2/3

Table 5. Final designation of importance to 100 × 100 m grid cells in Dresden (marked in shades of red): low (red 10%), moderate (red 35%), and high importance (red 55%). The designation of importance is based on detected sufficiency of parking spots (this table) and assigned relevance (Table 2).

N_{PB} \ P_{PB} [%]	0.1–3.0	3.1–8.0	8.1–18.0	18.1–35.0
1–3	/	/	/	/
4–8	Moderately insufficient to insufficient number of parking spots	Insufficient number of parking spots	/	/
9–18	Moderately insufficient number of parking spots	Insufficient number of parking spots	Moderately insufficient to insufficient number of parking spots	/
19–35	/	/	Moderately insufficient to insufficient number of parking spots	Moderately insufficient to insufficient number of parking spots
36–65	/	Moderately insufficient to insufficient number of parking spots	Moderately insufficient number of parking spots	Insufficient number of parking spots

5.3. Summary of Results

Based on the information gained from the previous steps, we were able to classify urban areas according to the parking sufficiency into areas with *moderately insufficient*, *moderately insufficient to insufficient* or *insufficient* number of parking spots. With this approach, we are able to identify the most critical areas in Dresden relating to bicycle parking. In Table 5 we present the results for each combination of N_{PB} and P_{PB} : for each cell, we report the sufficiency of parking spots and mark in shades of red whether the area finally has a *low*, *moderate*, or *high importance*. By comparing Tables 3 and 4, we assess that the areas around the central train station and the center of Altstadt are the most critical in Dresden. Foremost, within these locations, there is less than 50% of Category I cells of the overall cell number, while the number of Category III cells exceeds 50% for five

out of six categories. Therefore, we identify that these cells possess an insufficient or moderately insufficient to insufficient number of parking spots. Second, we consider the Flickr data related to these locations relevant because a significant percentage of posted photos contains detections of parked bicycles (up to 35%). Third, there is a high number of parked bicycle detections in these photos (up to 65). Finally, we previously classified those cells as cells with high relevance, which adds to the importance of results. We conclude that these locations qualify as a priority to be further inspected by urban planners in Dresden using other available data sources. Following the same approach, it is equally possible to classify each grid cell for a more detailed overview, which we skip for this paper.

Alongside, we identify that some areas also appear critical according to the number of bicycle parking spots, but due to the low percentage of photos containing parked bicycle detection, we classify them into low relevance cells. Subsequently, we consider them much less critical, and, finally, assign them low importance. Needless to say, this does not mean that those cells cannot be inspected further after the more critical locations have been resolved.

Our results also show that some cells contain significantly more than enough parking capacity that it appears to be in demand (Figure 14). We detect six cells of that type and suggest them as locations that could be further inspected to gain insights that could prove useful for improving future decisions related to the planning of the bicycle parking.

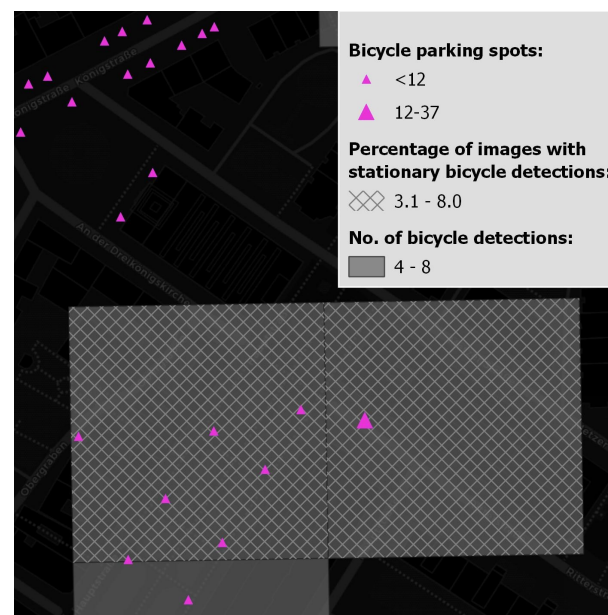


Figure 14. Example of a visually identified 100 × 100 m grid cell (left) that contains significantly more than enough bicycle parking capacity that it appears to be in demand. There are seven parking areas of capacity up to 12 bicycles, while there are only 4 to 8 bicycles detected inside the cell (N_{PB}).

We also observed that social media data provide more data for some areas and barely any for some others. That indicates that more popular locations offer more relevant bicycle-related data, while unpopular areas remain poorly covered. This popularity relates to the popularity within the used Flickr dataset. For example, the part of the Regierungsviertel district that is surrounded by the streets Albertstraße, Wigardstraße, and Glacisstraße demonstrates a significant number of available bicycle parking spots; however, we did not detect any parked bicycles there. Considering the number of parking spots, we conclude that the area is regularly frequented by cyclists but not interesting enough for posting it on Flickr. The bicycle parking racks there are installed around a car parking area, several residential buildings, schools, and institutions, such as the Saxon State Ministry of Justice. The same effect is also visible, e.g., along the Holbeinstraße and Tatzberg streets in the Johannstadt district where the racks are installed in front of residential buildings, an

athletic facility, research institutes, and a city's communal service company. According to a significant number of installed racks, it is apparent that these locations are also well frequented by cyclists but rarely posted on social media, such as Flickr.

In this chapter, we presented an analysis of the urban space of Dresden. We demonstrated the way social media data can be used to gain information related to the bicycle parking situation in an urban area. Additionally, we detected urban areas that, in our opinion, need prioritized attention from urban planning experts in the sense of further analysis regarding potentially missing bicycle parking. In the following chapter, we discuss the results of our research and introduce them in the context of their usefulness for the described task in urban planning.

6. Discussion

6.1. Object Detection on Social Media Data as an Additional Data Source for Obtaining Bicycle-Related Information in Urban Areas

We introduced a new method of obtaining bicycle-related data in urban areas and demonstrated it in the city of Dresden. In the case study, we showed that our method provides valuable information for planning urban infrastructure as we were able to identify areas with a lack of parking facilities for bicycles. Since our approach is focused on images from social media instead of commercial data sources or surveys, we are able to cover both individually owned and publicly shared bicycles within a geographical spread over the city that is typical for social media data [45]. In Section 4.4, we briefly compare the distribution of detected bicycles to other data sources and describe the dockless bicycle-sharing system (BSS) data as a biased in-the-wild sensing of the same area as our method. Regarding the location, some of the areas that we identified as having a moderately insufficient to insufficient number of parking facilities are also visible as dense bicycle clusters in the data of MOBI, such as the central train station. Up to this point, our method can be seen as a substitute for collecting data from bicycle sharing systems. This can already be interesting for cities without a BSS or companies planning to expand their BSS to a new city as a source for modeling a bicycle network and station location finding [46].

Compared to a BSS, we can provide additional value for urban planners for areas that have a higher frequency of bicycle detections but are not sufficiently covered by the BSS service. This often relates to the most scenic areas where shared bicycles are not allowed to be dropped off. In Dresden, for example, we identified central areas in the district Altstadt to have a moderately insufficient to insufficient number of bicycle parking facilities, but returning BSS bicycles was not allowed there. In some areas in Dresden that are also not covered by the return zones of the local BSS company, we did not identify lacking parking facilities: these are, e.g., the city park Großer Garten and the cycling route along the Elbe river. There were areas that we identified in our case study as not interesting enough for posting them on Flickr (i.e., social media). We consider this to be a rather minor drawback compared to BSS data because only some areas were frequently visible in the MOBI dataset (e.g., Johannstadt district), while there were no BSS bicycles present in other areas (e.g., in the Regierungsviertel district).

Additionally, we used a Flickr dataset that contains data from 2004 to 2014, but the popularity of social media has kept growing ever since, considering the fact that the number of mobile subscriptions increased from 2.33 billion in 2014 to 6.4 billion in 2021 [47]. We thus argue that our method can definitely provide additional value to urban planners, especially by using the most recent data and from more popular social media platforms that also rely on posting images, such as Instagram [48,49]. In that case, it would be necessary to additionally preprocess the data to comply with the privacy regulations by anonymizing them [50], e.g., as recently proposed by [51,52].

Regarding the efficiency of the bicycle detection, our approach achieves a recall of 81.4% (cf. Section 4.1) and thus does not find every bicycle, but it is significantly more efficient than manual identification by humans. We found that, on average, careful manual annotation (cf. Section 3.1) takes roughly 30s per image. While such manual detections may provide a significantly higher recall than our fully automatic approach, it would be

very costly to apply at a large-scale (e.g., thousands of images) setting. Our method, on the other hand, only requires a moderately powerful computer with comparatively small running costs.

In general, data availability is the only limiting factor for transferability in our approach, emphasizing the opportunities social media data can bring to urban analysis [53]. Requirements regarding computing capacity are rather low for applying the method to other cities, as the pre-trained object detection algorithm is utilized for identifying and classifying bicycles, and conventional GIS operations enable spatial exploration, analysis and visualization [54]. In comparison to the other methods for obtaining bicycle traffic data mentioned in Section 2, our approach is completely feasible without any structural installation, e.g., stationary counting stations or bicycle-mounted GPS sensors.

6.2. Relevance of Time in Our Approach

As we used the YFCC100m dataset, which covers a time span of 10 years, we only worked with temporally compressed location data. Compared to other bicycle-related data sources, this is at first a major drawback of our approach. Nevertheless, different spatio-temporal analyses would still be feasible: First, it is possible to analyze cumulated frequencies of bicycle detections in the dataset with respect to temporal categories such as weekdays, months or years [55]. These spatial-temporal patterns can provide insights into daily and seasonal routines of urban cyclists, and complement and verify the continuous data streams of stationary counters and BSS data [25,56]. Second, our method could be used as an indicator for the success of existing infrastructure, such as parking facilities or bicycle highways. It would be possible to compare the spatial pattern of detected bicycles before and after the date of construction and, therefore, detect changes in bicycle-related traffic flows or the density of detected bicycles for each grid cell we used in our case study. The capability of social media to acquire information about interventions in different events has already been shown in numerous studies, even though most of them examined more significant interventions such as in cases of traffic incidents or natural disasters [57,58].

6.3. Influence of Bicycle Detection Errors and Potential Improvements

In Section 4.1, we present the results of our object detection methodology for identifying moving and stationary bicycles. In addition to a quantitative evaluation, we also provide examples of failure cases and identify the most frequent categories of false negatives and false positives. Although we conclude—based on our case study—that the information extracted from social media data is suitable for identifying a lack of parking facilities for bicycles, we are aware of the flaws of our object detection pipeline. Further on, we discuss their impact on the findings of the case study as well as potential improvements within our approach.

With regards to content, the detection errors can be subdivided into strictly bicycle-related errors (i.e., missed bicycles or duplicate detections) and those of confusion with other objects such as traffic signs or chairs. For the former, we argue that the impact of the false detections on the findings of the case study is relatively small, as we tuned the algorithm to operate using more conservative thresholds and, therefore, we rather tend to underestimate the number of bicycles. For images with multiple bicycles detected, this will decrease the total number of bicycles, but the number of images with bicycle detections will not be affected. With Tables 2 and 3 in mind, changing a class in the vertical direction of the table will most likely not change the assigned relevance of the grid cell. On the other hand, a false negative (i.e., missed detection) of a lone bicycle in an image can change the relevance according to Table 2, but according to the distribution of the classes in Table 3, a relevance change will most likely occur only between classes of low and medium relevance. If we overestimate the number of bicycles in certain images, or occasionally detect false positives of *other wheeled objects* such as motorcycles, we argue that assigning a higher relevance to these cells is an acceptable drawback for the focus of our case study, where we qualify locations as a priority to be further inspected by urban planners.

In contrast to bicycle-related detections, to improve our approach, we consider minimizing the false positive detections of miscellaneous objects such as traffic signs. As we mention in Section 3.2, we use an object detector that has been trained on the very diverse COCO [40] dataset, which contains a large number of object classes in images taken across all continents. Our target application, however, is narrower: we are only interested in the detection of bicycles and persons, and in our case study, the environment is limited to the city of Dresden. In such a case, fine-tuning the object detection neural network on the narrower target domain has been shown to improve performance [59]. In addition, it can be feasible to combine our method with other detection algorithms from computer vision. For example, one can employ an algorithm that detects bicycle traffic signs [60,61] in order to filter out this category of false positive detections. Similarly, one could employ object detection or image classification [62] algorithms trained on other classes than bicycle and person in order to estimate the likelihood of a bicycle detection being a false positive due to confusion. Bicycle detections, both correct and false, may also occur in images taken indoors. As these are of no interest for our application, it would be reasonable to discard such images automatically using an algorithm that can distinguish indoor from outdoor scenes [63]. For improving the *moving vs. stationary* classification accuracy, it would be feasible to train an object detection algorithm to directly provide this classification instead of relying on person detections as we describe in Section 3.2.1. Whether this would actually work better, however, depends on a variety of factors, such as the size and diversity of the training dataset, the architecture of the convolutional neural network, and appropriate data augmentation during training [64].

Improving bicycle detection from social media posts might also include processing tags, textual descriptions, and emoticons used in a post to extract user reactions [65]. Sentiment analysis and categorization of emotions associated with a post can be applied to identify if postings related to bicycles are more positive or negative connoted and in that way obtain detailed contextual information related to a post and specific areas [66,67]. Learning about the context of an image could allow to, e.g., eliminate false detections, such as musical instruments and tripods on concert stages.

If an image contains bicycles that are located at a far distance, using only the GPS location associated with the image may result in poor localization of the bicycles. It would be feasible, however, to use visual localization [68] and geometric cues [69] in order to estimate the pose of the camera and, subsequently, extract 3D information about the scene from the image [70–72]. This would provide a more precise localization of the detected bicycles.

7. Conclusions

In this paper, we introduced a new method of obtaining bicycle-related data from social media posts. In the first step, we used a pre-trained state-of-the-art object detection algorithm to detect bicycles on a regional subset of the YFCC100m dataset. In the second step, we differentiated between moving and stationary bicycles as we leveraged the ability of the detection algorithm to detect people and assume that if a bicycle is located right next to or below a person, it is non-stationary or moving. With our method, we detected 4157 bicycles in the city of Dresden with an overall precision of 96.1% and a recall of 81.4%, and classified 85.8% of the moving bicycles and 91.9% of the stationary bicycles correctly. We then conducted a case study, where we analyzed the general situation in Dresden related to parking facilities for bicycles using the results of the object detection. As a result, we were able to classify urban areas according to the sufficiency of parking facilities and thereby identify areas that need prioritized attention from urban planners. Using the same approach, it would be possible to detect further micro-locations within these urban areas.

Our method proved as relevant because we were able to gain meaningful insights into the urban area of Dresden using social media data. We conclude that it provides significant value in planning the bicycle infrastructure in a city, particularly considering that data for that purpose is otherwise difficult and expensive to collect. We were able to shorten the

data acquisition multiple times in comparison to traditional methods of data collection in urban planning. Additionally, we can also provide a much larger temporal coverage.

Clear limitations of using social media data lie in the fact that the data availability and coverage largely depend on the usage of social media. The spatial coverage is worse for unpopular than for popular urban areas, and temporal coverage (i.e., days, months, years, etc.) is worse for years when social media was less popular. However, both data availability and coverage depend on usage trends of the social media platform whose data were used. Having that in mind, usage of the newest social media data should enable even more precise spatio-temporal analyses. This precision may additionally be increased by implementing further object detection algorithms in data processing.

By choosing Dresden as our case study, we benefited from a manageable amount of data, as well as from our own local expertise. In the future, we intend to potentially enlarge the amount of bicycle detections to deal with by focusing on a larger city. Another way to enlarge our dataset would be to integrate data from more social media platforms such as Instagram or Twitter, depending on the availability of the data. We also consider implementing other databases of interest, especially those providing street-level images such as Mapillary. Furthermore, dealing with a higher number of detections and larger urban areas of interest requires more sophisticated approaches of localization and visualization, so we intend to work on methods to orientate the images by matching detected objects and city furniture (e.g., benches, lightning objects, etc.) and improve the techniques for visual analysis and exploration.

Author Contributions: Conceptualization, Moris Zahtila, Florian Kluger and Martin Knura; methodology, Moris Zahtila, Florian Kluger and Martin Knura; software, Florian Kluger; validation, Moris Zahtila, Florian Kluger and Martin Knura; formal analysis, Moris Zahtila, Florian Kluger and Martin Knura; investigation, Moris Zahtila, Florian Kluger and Martin Knura; resources, Bodo Rosenhahn, Jochen Schiewe and Dirk Burghardt; data curation, Florian Kluger and Martin Knura; writing—original draft preparation, Moris Zahtila, Florian Kluger and Martin Knura; writing—review and editing, Moris Zahtila, Florian Kluger, Martin Knura, Bodo Rosenhahn, Jochen Schiewe and Dirk Burghardt; visualization, Moris Zahtila and Martin Knura; supervision, Bodo Rosenhahn, Jochen Schiewe and Dirk Burghardt; project administration, Moris Zahtila, Florian Kluger, Martin Knura, Bodo Rosenhahn, Jochen Schiewe and Dirk Burghardt; funding acquisition, Bodo Rosenhahn, Jochen Schiewe and Dirk Burghardt. All authors have read and agreed to the published version of the manuscript.

Funding: This collaboration was realized within the DFG Priority Programme (SPP 1894/2) and supported by grants COVMAP (RO 2497/12-2), TOVIP (SCHI 1008/11-1) and EVA-VGI 2 (BU 2605/8-2).

Data Availability Statement: Data and source code is available at https://github.com/fkluger/bicycle_detection.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Sumantran, V.; Fine, C.; Gonsalvez, D. *Faster, Smarter, Greener: The Future of the Car and Urban Mobility*; The MIT Press: Cambridge, MA, USA, 2017. [\[CrossRef\]](#)
2. Qiu, L.Y.; He, L.Y. Bike Sharing and the Economy, the Environment, and Health-Related Externalities. *Sustainability* **2018**, *10*, 1145. [\[CrossRef\]](#)
3. Pucher, J.; Buehler, R. Making Cycling Irresistible: Lessons from The Netherlands, Denmark and Germany. *Transp. Rev.* **2008**, *28*, 495–528. [\[CrossRef\]](#)
4. Froehlich, J.; Neumann, J.; Oliver, N. Sensing and Predicting the Pulse of the City through Shared Bicycling. In Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, CA, USA, 14–17 July 2009; pp. 1420–1426.
5. Zhou, X. Understanding Spatiotemporal Patterns of Biking Behavior by Analyzing Massive Bike Sharing Data in Chicago. *PLoS ONE* **2015**, *10*, e0137922. [\[CrossRef\]](#)
6. Corcoran, J.; Li, T.; Rohde, D.; Charles-Edwards, E.; Mateo-Babiano, D. Spatio-temporal patterns of a Public Bicycle Sharing Program: The effect of weather and calendar events. *J. Transp. Geogr.* **2014**, *41*, 292–305. [\[CrossRef\]](#)

7. Etienne, C.; Latifa, O. Model-Based Count Series Clustering for Bike Sharing System Usage Mining: A Case Study with the VéLib' System of Paris. *ACM Trans. Intell. Syst. Technol.* **2014**, *5*, 1–21. [[CrossRef](#)]
8. Korpilo, S.; Virtanen, T.; Lehvävirta, S. Smartphone GPS tracking—Inexpensive and efficient data collection on recreational movement. *Landsc. Urban Plan.* **2017**, *157*, 608–617. [[CrossRef](#)]
9. Reades, J.; Calabrese, F.; Sevtsuk, A.; Ratti, C. Cellular Census: Explorations in Urban Data Collection. *IEEE Pervasive Comput.* **2007**, *6*, 30–38. [[CrossRef](#)]
10. Huber, S.; Lissner, S.; Schnabel, A.; Lindemann, P.; Friedl, J. Modelling bicycle route choice in German cities using open data, MNL and the bikeSim web-app. In Proceedings of the 2021 7th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS), Heraklion, Greece, 16–17 June 2021. [[CrossRef](#)]
11. Chen, C.; Anderson, J.C.; Wang, H.; Wang, Y.; Vogt, R.; Hernandez, S. How bicycle level of traffic stress correlate with reported cyclist accidents injury severities: A geospatial and mixed logit analysis. *Accid. Anal. Prev.* **2017**, *108*, 234–244. [[CrossRef](#)] [[PubMed](#)]
12. Götschi, T.; Garrard, J.; Giles-Corti, B. Cycling as a Part of Daily Life: A Review of Health Perspectives. *Transp. Rev.* **2016**, *36*, 45–71. [[CrossRef](#)]
13. Buehler, R. Determinants of transport mode choice: A comparison of Germany and the USA. *J. Transp. Geogr.* **2011**, *19*, 644–657. [[CrossRef](#)]
14. Scott, D.M.; Lu, W.; Brown, M.J. Route choice of bike share users: Leveraging GPS data to derive choice sets. *J. Transp. Geogr.* **2021**, *90*, 102903. [[CrossRef](#)]
15. Hoogendoorn, S.; Daamen, W. Bicycle Headway Modeling and Its Applications. *Transp. Res. Rec.* **2016**, *2587*, 34–40. [[CrossRef](#)]
16. Pogodzinska, S.; Kiec, M.; D'Agostino, C. Bicycle Traffic Volume Estimation Based on GPS Data. *Transp. Res. Procedia* **2020**, *45*, 874–881.
17. Gehrke, S.R.; Reardon, T.G. Direct demand modelling approach to forecast cycling activity for a proposed bike facility. *Transp. Plan. Technol.* **2021**, *44*, 1–15. [[CrossRef](#)]
18. Beitel, D.; McNee, S.; Miranda-Moreno, L.F. Quality Measure of Short-Duration Bicycle Counts. *Transp. Res. Rec.* **2017**, *2644*, 64–71. [[CrossRef](#)]
19. Turner, S.; Lasley, P. Quality Counts for Pedestrians and Bicyclists: Quality Assurance Procedures for Nonmotorized Traffic Count Data. *Transp. Res. Rec.* **2013**, *2339*, 57–67. [[CrossRef](#)]
20. Hankey, S.; Lu, T.; Mondschein, A.; Buehler, R. Spatial models of active travel in small communities: Merging the goals of traffic monitoring and direct-demand modeling. *J. Transp. Health* **2017**, *7*, 149–159. [[CrossRef](#)]
21. Gupte, S.; Masoud, O.; Martin, R.; Papanikolopoulos, N. Detection and classification of vehicles. *IEEE Trans. Intell. Transp. Syst.* **2002**, *3*, 37–47. [[CrossRef](#)]
22. Ghosh, A.; Sabuj, M.S.; Sonet, H.H.; Shatabda, S.; Farid, D.M. An Adaptive Video-based Vehicle Detection, Classification, Counting, and Speed-measurement System for Real-time Traffic Data Collection. In Proceedings of the 2019 IEEE Region 10 Symposium (TENSYMP), Kolkata, India, 7–9 June 2019; pp. 541–546. [[CrossRef](#)]
23. Gillis, D.; Gautama, S.; Van Gheluwe, C.; Semanjski, I.; Lopez, A.J.; Lauwers, D. Measuring Delays for Bicycles at Signalized Intersections Using Smartphone GPS Tracking Data. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 174. [[CrossRef](#)]
24. Chen, Z.; van Lierop, D.; Ettema, D. Dockless bike-sharing systems: What are the implications? *Transp. Rev.* **2020**, *40*, 333–353. [[CrossRef](#)]
25. Ma, X.; Ji, Y.; Yuan, Y.; Van Oort, N.; Jin, Y.; Hoogendoorn, S. A comparison in travel patterns and determinants of user demand between docked and dockless bike-sharing systems using multi-sourced data. *Transp. Res. Part A Policy Pract.* **2020**, *139*, 148–173. [[CrossRef](#)]
26. Basiri, A.; Haklay, M.; Foody, G.; Mooney, P. Crowdsourced geospatial data quality: Challenges and future directions. *Int. J. Geogr. Inf. Sci.* **2019**, *33*, 1588–1593. [[CrossRef](#)]
27. Wang, L. Planning for cycling in a growing megacity: Exploring planners' perceptions and shared values. *Cities* **2020**, *106*, 102857. [[CrossRef](#)]
28. Weng, J.; Bäumer, T.; Müller, P. Bike-Sharing Systems as Integral Components of Inner-City Mobility Concepts: An Analysis of the Intended User Behaviour of Potential and Actual Bike-Sharing Users. In *Innovations for Metropolitan Areas: Intelligent Solutions for Mobility, Logistics and Infrastructure Designed for Citizens*; Planing, P., Müller, P., Dehdari, P., Bäumer, T., Eds.; Springer: Berlin/Heidelberg, Germany, 2020; pp. 121–132. [[CrossRef](#)]
29. Zhu, Z.; Blanke, U.; Calatroni, A.; Tröster, G. Human Activity Recognition Using Social Media Data. In Proceedings of the 12th International Conference on Mobile and Ubiquitous Multimedia, Luleå, Sweden, 2–5 December 2013. [[CrossRef](#)]
30. Stock, K. Mining location from social media: A systematic review. *Comput. Environ. Urban Syst.* **2018**, *71*, 209–240. [[CrossRef](#)]
31. Middleton, S.E.; Kordopatis-Zilos, G.; Papadopoulos, S.; Kompatsiaris, Y. Location Extraction from Social Media: Geoparsing, Location Disambiguation, and Geotagging. *ACM Trans. Inf. Syst.* **2018**, *36*, 1–27. [[CrossRef](#)]
32. Gong, J.; Li, R.; Yao, H.; Kang, X.; Li, S. Recognizing Human Daily Activity Using Social Media Sensors and Deep Learning. *Int. J. Environ. Res. Public Health* **2019**, *16*, 3955. [[CrossRef](#)]
33. Cao, R.; Tu, W.; Yang, C.; Li, Q.; Liu, J.; Zhu, J.; Zhang, Q.; Li, Q.; Qiu, G. Deep learning-based remote and social sensing data fusion for urban region function recognition. *ISPRS J. Photogramm. Remote Sens.* **2020**, *163*, 82–97. [[CrossRef](#)]

34. Thomee, B.; Shamma, D.A.; Friedland, G.; Elizalde, B.; Ni, K.; Poland, D.; Borth, D.; Li, L.J. YFCC100M: The new data in multimedia research. *Commun. ACM* **2016**, *59*, 64–73. [CrossRef]
35. Tan, M.; Pang, R.; Le, Q.V. Efficientdet: Scalable and efficient object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–18 June 2020; pp. 10781–10790.
36. Redmon, J.; Farhadi, A. Yolov3: An incremental improvement. *arXiv* **2018**, arXiv:1804.02767.
37. Du, X.; Lin, T.Y.; Jin, P.; Ghiasi, G.; Tan, M.; Cui, Y.; Le, Q.V.; Song, X. SpineNet: Learning scale-permuted backbone for recognition and localization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–18 June 2020; pp. 11592–11601.
38. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask r-cnn. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2961–2969.
39. Lin, T.Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. Focal loss for dense object detection. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2980–2988.
40. Lin, T.Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; Zitnick, C.L. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*; Springer: Berlin/Heidelberg, Germany, 2014; pp. 740–755.
41. Zoph, B.; Cubuk, E.D.; Ghiasi, G.; Lin, T.Y.; Shlens, J.; Le, Q.V. Learning data augmentation strategies for object detection. In *European Conference on Computer Vision*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 566–583.
42. Kuhn, H.W. The Hungarian method for the assignment problem. *Nav. Res. Logist. Q.* **1955**, *2*, 83–97. [CrossRef]
43. Landeshauptstadt Dresden, S.u.T. Dauerzählstellen für den Radverkehr. Available online: http://www.dresden.de/media/pdf/Strassenbau/Dauerzaehlstellen_Stadtplan.pdf (accessed on 5 August 2021).
44. Jenks, G.F. The data model concept in statistical mapping. *Int. Yearb. Cartogr.* **1967**, *7*, 186–190.
45. Jiang, B.; Ma, D.; Yin, J.; Sandberg, M. Spatial Distribution of City Tweets and Their Densities. *Geogr. Anal.* **2016**, *48*, 337–351. [CrossRef]
46. Bahadori, M.S.; Gonçalves, A.B.; Moura, F. A Systematic Review of Station Location Techniques for Bicycle-Sharing Systems Planning and Operation. *ISPRS Int. J. -Geo-Inf.* **2021**, *10*, 554. [CrossRef]
47. Telefonaktiebolaget LM Ericsson Ericsson Mobility Visualizer. Available online: <https://www.ericsson.com/en/mobility-report/mobility-visualizer?f=1&ft=1&r=1&t=8&s=1&u=1&y=2014,2021&c=1> (accessed on 14 September 2021).
48. Domínguez, D.R.; Díaz Redondo, R.P.; Vilas, A.F.; Khalifa, M.B. Sensing the city with Instagram: Clustering geolocated data for outlier detection. *Expert Syst. Appl.* **2017**, *78*, 319–333. [CrossRef]
49. Gunter, U.; Önder, I. An Exploratory Analysis of Geotagged Photos From Instagram for Residents of and Visitors to Vienna. *J. Hosp. Tour. Res.* **2021**, *45*, 373–398. [CrossRef]
50. Beigi, G.; Shu, K.; Zhang, Y.; Liu, H. Securing Social Media User Data: An Adversarial Approach. In Proceedings of the 29th on Hypertext and Social Media, Baltimore, MD, USA, 9–12 July 2018; pp. 165–173. [CrossRef]
51. Dunkel, A.; Löchner, M.; Burghardt, D. Privacy-Aware Visualization of Volunteered Geographic Information (VGI) to Analyze Spatial Activity: A Benchmark Implementation. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 607. [CrossRef]
52. Löchner, M.; Fathi, R.; Schmid, D.; Dunkel, A.; Burghardt, D.; Fiedrich, F.; Koch, S. Case Study on Privacy-Aware Social Media Data Processing in Disaster Management. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 709. [CrossRef]
53. Martí, P.; Serrano-Estrada, L.; Nolasco-Cirugeda, A. Social Media data: Challenges, opportunities and limitations in urban studies. *Comput. Environ. Urban Syst.* **2019**, *74*, 161–174. [CrossRef]
54. Alivand, M.; Hochmair, H.H. Spatiotemporal analysis of photo contribution patterns to Panoramio and Flickr. *Cartogr. Geogr. Inf. Sci.* **2017**, *44*, 170–184. [CrossRef]
55. McKenzie, G.; Janowicz, K.; Gao, S.; Yang, J.A.; Hu, Y. POI Pulse: A Multi-granular, Semantic Signature-Based Information Observatory for the Interactive Visualization of Big Geosocial Data. *Cartographica* **2015**, *50*, 71–85. [CrossRef]
56. Nordback, K.; Marshall, W.E.; Janson, B.N.; Stolz, E. Estimating Annual Average Daily Bicyclists: Error and Accuracy. *Transp. Res. Rec.* **2013**, *2339*, 90–97. [CrossRef]
57. Rashidi, T.H.; Abbasi, A.; Maghrebi, M.; Hasan, S.; Waller, T.S. Exploring the capacity of social media data for modelling travel behaviour: Opportunities and challenges. *Transp. Res. Part C Emerg. Technol.* **2017**, *75*, 197–211. [CrossRef]
58. Resch, B.; Usländer, F.; Havas, C. Combining machine-learning topic models and spatiotemporal analysis of social media data for disaster footprint and damage assessment. *Cartogr. Geogr. Inf. Sci.* **2018**, *45*, 362–376. [CrossRef]
59. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 580–587.
60. Reinders, C.; Ackermann, H.; Yang, M.Y.; Rosenhahn, B. Object recognition from very few training examples for enhancing bicycle maps. In Proceedings of the 2018 IEEE Intelligent Vehicles Symposium (IV), Suzhou, China, 26–30 June 2018; pp. 1–8.
61. Reinders, C.; Ackermann, H.; Yang, M.Y.; Rosenhahn, B. Learning convolutional neural networks for object detection with very little training data. In *Multimodal Scene Understanding*; Elsevier: Amsterdam, The Netherlands, 2019; pp. 65–100.
62. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **2012**, *25*, 1097–1105. [CrossRef]
63. Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; Torralba, A. Places: A 10 million Image Database for Scene Recognition. *IEEE Transactions Pattern Anal. Mach. Intell.* **2017**, *40*, 1452–1464. [CrossRef]

64. Kluger, F.; Reinders, C.; Raetz, K.; Schelske, P.; Wandt, B.; Ackermann, H.; Rosenhahn, B. Region-based cycle-consistent data augmentation for object detection. In Proceedings of the 2018 IEEE International Conference on Big Data (Big Data), Seattle, WA, USA, 10–13 December 2018; pp. 5205–5211.
65. Dunkel, A.; Andrienko, G.; Andrienko, N.; Burghardt, D.; Hauthal, E.; Purves, R. A conceptual framework for studying collective reactions to events in location-based social media. *Int. J. Geogr. Inf. Sci.* **2019**, *33*, 780–804. [[CrossRef](#)]
66. Hauthal, E.; Burghardt, D. Mapping Space-Related Emotions out of User-Generated Photo Metadata Considering Grammatical Issues. *Cartogr. J.* **2016**, *53*, 78–90. [[CrossRef](#)]
67. Hauthal, E.; Burghardt, D.; Dunkel, A. Analyzing and Visualizing Emotional Reactions Expressed by Emojis in Location-Based Social Media. *ISPRS Int. J. Geo-Inf.* **2019**, *8*, 113. [[CrossRef](#)]
68. Sarlin, P.E.; Cadena, C.; Siegwart, R.; Dymczyk, M. From coarse to fine: Robust hierarchical localization at large scale. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019; pp. 12716–12725.
69. Kluger, F.; Ackermann, H.; Yang, M.Y.; Rosenhahn, B. Temporally consistent horizon lines. In Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May–31 August 2020; pp. 3161–3167.
70. Kluger, F.; Ackermann, H.; Yang, M.Y.; Rosenhahn, B. Deep learning for vanishing point detection using an inverse gnomonic projection. In *German Conference on Pattern Recognition*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 17–28.
71. Kluger, F.; Brachmann, E.; Ackermann, H.; Rother, C.; Yang, M.Y.; Rosenhahn, B. Consac: Robust multi-model fitting by conditional sample consensus. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 4634–4643.
72. Kluger, F.; Ackermann, H.; Brachmann, E.; Yang, M.Y.; Rosenhahn, B. Cuboids Revisited: Learning Robust 3D Shape Fitting to Single RGB Images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 19–25 June 2021; pp. 13070–13079.