Descriptive route choice analysis of cyclists in Zurich

Adrian Meister

Isha Gupta

Kay W. Axhausen

STRC | 21st Swiss Transport Research Conference
Monte Verità / Ascona, September 12 – 14, 2021

Institute for Transport Planning and Systems

Descriptive route choice analysis of cyclists in Zurich

Adrian Meister
Institute for Transport Planning
and Systems (IVT)
ETH Zurich
CH-8093
adrian.meister@ivt.baug.ethz.ch

Isha Gupta
Institute for Transport Planning
and Systems (IVT)
ETH Zurich
CH-8093
igupta@student.ethz.ch

Kay W. Axhausen
Institute for Transport Planning
and Systems (IVT)
ETH Zurich
CH-8093
axhausen@ivt.baug.ethz.ch

September 2021

Abstract

This paper presents (preliminary) results of a descriptive route choice analysis of cyclists in Zurich. Raw GPS trajectories from the MOBIS-COVID dataset are matched to an enriched OSM network. Two different map-matching algorithms are evaluated based on their sensitivity to hyper-parameters as well as different network configurations. The results of this paper lay the ground work for future modelling of the data using discrete route choice frameworks.

Keywords

Cycling, cycling behavior, route choice, map matching

1 Introduction

Cycling is becoming an increasingly popular mode of transport in many regions of the world. Especially the COVID-19 pandemic has generated a surge in bicycles sales and cycling activity around the globe. The probably most important advantage bicycles have from a system-wide perspective is their very low carbon footprint. Considering the life-cycle emissions of different urban transport modes for the current global average electricity mix, personal bicycles and e-bicycles produce by far the lowest amount of CO_2-equivalent per person/km (Blondel et al., 2011). Even under a clean-energy assumption, bicycles and e-bicycles clearly outperform motorized individual and micro-mobility modes of transport (Cazzola and Crist, 2020). Cycling can therefore act as a useful contributor to the decarbonization of transport. Furthermore, cycling also contributes to solve urban congestion, which are assumed to become even more severe due to the expected growth of population and urban areas. Compared to conventional motorized traffic like cars, bicycles are efficient in terms of their required space, and are of advantage when designing transport system on constrained urban networks with limited capacity (Sims et al., 2014). The climate crisis and urbanization (and even pandemics) are long-term drivers at a global scale, which we can assume to further drive the demand for cycling in the future. Considering the typical distance of urban motorized trips shows the potential for modal shifts in favor of cycling. Data from the latest Swiss mobility micro-census (MZMV) (BFS, 2017) indicates that around 50% of urban car trips are less than 5km and around 65% are less than 10km. These are equivalent to approx. 20 mins of cycling and electric cycling respectively. Increased shares of cycling however also come with limitations and drawbacks. On the one hand, cycling is heavily influenced by local climatic and geographic conditions, making it an inconvenient option for regions with certain terrain and precipitation characteristics. On the other hand, the increase of cycling in the last decade correlates with increasing accident numbers.

The growing importance of cycling for the development of sustainable and safe urban transport calls for appropriate modelling tools which can be used for transport planning. This paper presents the results of a descriptive analysis of the route choice of cycling trips from the MOBIS-COVID study (Molloy et al., 2021). It lays the ground work required to estimate up-to-date route choice models based on discrete choice frameworks, and the respective integration of such models into agent-based simulation tools like MATSim (Axhausen et al., 2016). The raw GPS data of the MOBIS-COVID dataset from bicycle trips is processed and enriched with empirically relevant and use-case specific attributes of interest. To do so, we evaluate two different Hidden-Markov-model (HMM) based map-matching pipelines. Both are optimized regarding their hyper-parameters and the underlying network to which the trajectories are matched to. We evaluate the outcomes on a quantitative and qualitative basis and investigate the resulting enrichment.

2 Related Work

2.1 Cyclist route choice

Modelling cycling route choices has been subjects to numerous previous studies. While most of these studies are based on stated-preference (SP) data, an increasing amount is based on revealed-preference (RP) data in form of (GPS-enabled) travel diary surveys. SP data has the weakness that it is unclear to which degree the stated preferences correspond to actual choices made in a real environment (Broach et al., 2012). RP data has become easier to collect, especially through smartphones, and is seen as more robust (Casello and Usyukov, 2014). Studies that use GPS-tracking data to estimate cycling route choice models among others include Hood et al. (2011); Zimmermann et al. (2017) and Misra and Watkins (2018). Some of the latest published models like Lu et al. (2018) and Khatri et al. (2016) focus on bike-sharing systems. There is currently only the work of Dane et al. (2019) that specifically considered electric bicycles. The wide body of literature covering both SP and RP data, mostly have similar findings about which attributes (trip, person, environment) have substantial effects on cycling route choice, with only few studies showing contradictory findings. These attributes include: route length, slope and travel time, cycling infrastructure, continuity and intersections, surface and smoothness, parking and bus stops, as well as traffic volumes, parks and urban vegetation.

Notable work includes the analysis of Menghini et al. (2010) and de Freitas et al. (2019), which both studied route choice of cyclists in Zurich. The former used a similar methodology than in this paper, i.e. map-matched GPS trajectories, and a dataset size of approx. 2'500 private bicycle trips from 2004. The considered choice attributes include trip length, average absolute gradient, maximum gradient, share of bike path (no differentiation between bike lanes and physically separated bike paths) and number of traffic lights along the route. The network was sourced from non-OSM sources and included approx. 25'000 links and 8'500 nodes. For matching the trajectories, the authors used the rule-based map matching approach from Schuessler and Axhausen (2009). Compared to then, the technology for measuring and processing GPS data has seen dramatic advances in performance and accuracy. Furthermore, the trend of e-bike adoption and usage has gained significant momentum through the last decade, and new forms of shared bike mobility have been introduced. The work of Menghini et al. (2010) can hence be considered out-dated regarding the subject as well as the applied technology stack. de Freitas et al. (2019) used the national household travel survey (BFS, 2017) diary data from 2015, which includes approx. 500 manually reported bicycle and e-bicycle trips in Zurich for which the origin-destination pairs are provided. The actual paths were routed using the Graphhopper (GH) routing engine (GraphHopper, 2020b) on a network with approx. 30'000 links and 13'000

nodes. The choice attributes which were used for modelling include only topological ones, i.e. distance and gradients, as the work focuses on intermodal trip-making and also considers attributes relevant to other modes.

2.2 Map Matching

Map matching denotes the process of assigning positional data (mostly in form of GPS coordinates) to some geographic information about the infrastructure, most commonly road networks. Map matching can be conducted offline using a prerecorded sequence of coordinates, but also in real time. Common applications include navigation systems, which must determine which road the user might be on in order to calculate a further route. There have been various algorithmic approaches proposed to solve the map matching problem. In terms of the current state-of-the-art, the approaches can be divided into four categories, as outlined in Chao et al. (2020): Similarity models, which return the vertices/edges that are closest to the trajectory geometrically and/or topologically; Candidate-evolving models, which maintain a set of candidate roads during the matching and continually evolves it by adding new candidates and excluding irrelevant ones based on a fixed set of rules; Scoring models, for example naive weighting which selects road edges from a group of candidates for each trajectory segment which maximises the predefined scoring function; State transition models, which build weighted topological graphs that contain all possible routes that the traveler might choose - the vertices represent the possible states the traveler may be located while the edges represent the transitions between states at different times. For a detailed overview of map matching algorithms, see Jensen and Tradišauskas (2009).

The majority of applications and implemented open-source tools are based on state transition models, more specifically HMMs, from which two different implementations are evaluated in this work. Newson and Krumm (2009) is often referred to as the state-of-the-art implementation of the HMM approach. HMMs are effective as they account for measurement noise and the layout of the road network to avoid a naive distance-based matching of a point to a road. The path selection is based on a set of emission probabilities, which represent the likelihood that a measurement resulted from a given state. Transitions are favoured if their distance is around equal to the great circle between the measurements. Finally, the Viterbi algorithm (see Forney (1973)) employs dynamic programming to quickly find the path through the lattice of the HMM graph that maximises the measurement probability and transition probabilities to compute the best match. It also takes into consideration that the input might be degraded, which is an important aspect when using notoriously noisy GPS measurements.

3 Methodology

The goal of the developed method is to map the GPS trajectories to a cycling network, while ensuring two central requirements. The first is providing the most complete representation of real-world conditions and trajectories with the highest possible degree of accuracy and spatial resolution. The second is to generate a highly enriched dataset which can be used for subsequent choice analysis. Two different implementations of the HMM-based map-matching approach are optimized regarding their hyper-parameters and evaluated against each other using different mode-specific sub-networks (car, bike, walk). The quantitative evaluation is extended by visually evaluating the plausibility of the matching results on a qualitative basis. For the best performing combination of algorithm and network, we finally evaluate the resulting enrichment mainly based on the resulting share of network- vs. matched trajectory coverage of each attribute.

3.1 Data

The data used for this analysis includes raw GPS data of bicycle and e-bicycle trips from the MOBIS-COVID dataset. The MOBIS-COVID project (Molloy et al., 2021) initially started as mobility pricing experiment in Switzerland. In September 2019, approx. 5,400 people started an app-based GPS-tracking study in which each respondent was tracked for eight weeks. About 3,600 respondents completed the whole study period. As of March 2020, 1,600 former participants volunteered to reactivate their accounts in order to capture the COVID-19 pandemic-related effects. The project is on-going and had an average of approx. 1,000 active participants tracked throughout all 2020. The study uses the Catch-my-Day app from MotionTag GmbH. The app runs in the background and collects GPS and other device sensory data (mainly acceleration). The underlying SDK is designed to minimize battery consumption and transfers the collected data to a backend. The incoming raw data is processed into a travel diary format with legs for which each's mode is detected using a deep recurrent neural network with an accuracy of around 90%. The app allows the user to change and validate the detected mode, providing a wider range of modes including such which are not detectable, e.g., e-bicycles or bike-sharing. The detailed socio-demographics of the respective traveler, as well as purpose and weather information are available for each trip.

The dataset includes approx. 1.3 Mio. data-points which make up approx. 110,000 (e-)bicycle trips for all of Switzerland. These trips are filtered for the geographical city boundaries of Zurich, resulting in 8,342 remaining trips. These potentially include non-bicycle trips due to miss-classification, for which reason they are further filtered based on plausible speeds (average
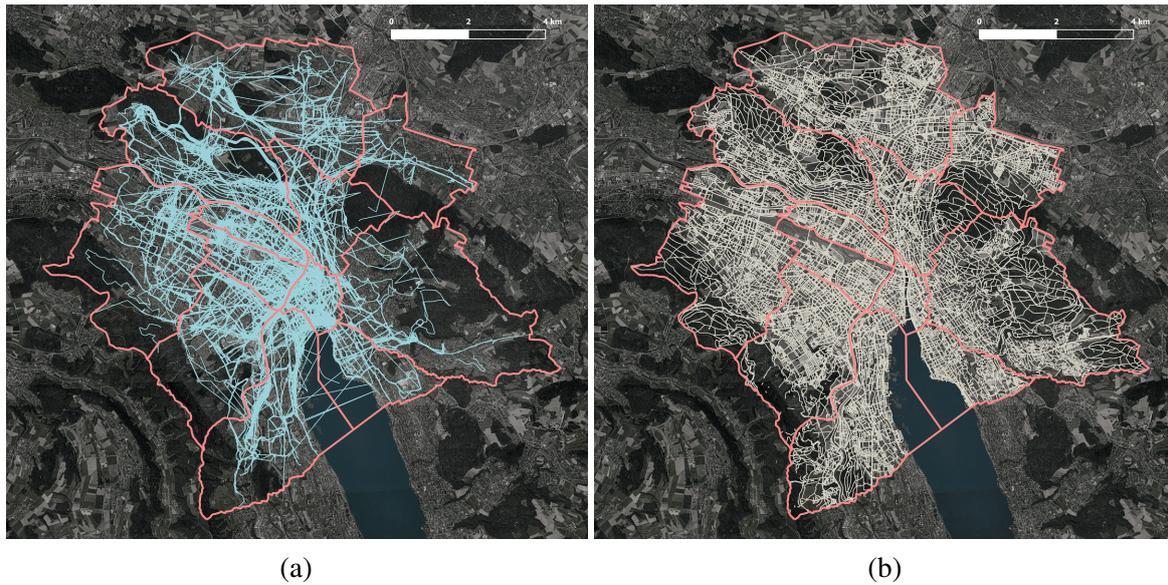
Figure 1: Study area (Zurich city districts boundaries in pink) with raw GPS trajectories (a, blue), and with OSM network (b, grey).

speed larger than 5km/h and smaller than 40km/h) and distances (GPS distance, i.e. summed point-wise distance of the trajectory of at least 1km). The resulting subset includes 5,656 matching candidates from 367 individuals, which are shown in Figure 1a. The average GPS distance is 2.3km, the average trip duration 9.2 minutes, and the average speed 16.8km/h. Further details about the trip purpose and traveler characteristics are listed in Table 5 and 6 in the Appendix.

3.2 Network

The network data is sourced from the OSM project. The complete OSM network for the city boundaries of Zurich is shown in Figure 1b. From this network, three mutually exclusive car-, bike-, and pedestrian layers are extracted based on relevant OSM-tags (mostly "highway" tag). The different layers are combined to directed graphs and used as basis for the sensitivity analysis later on. The car layer consists of all edges on which car-traffic is allowed, whereas "path" links are excluded. The layer is further cleaned for large highways on which cycling is prohibited or highly unlikely. The bike layer includes "path" and "cycleway" links, as well as all those which are not in the car layer, but for which cycling is specifically allowed (derived from the "bicycle" tag). The pedestrian layer includes the walkable parts of the OSM network, specifically using "pedestrian" and "footway" links. The characteristics of the considered networks are listed in Table 1. One can see that the bike and pedestrian layer add a considerable amount of nodes and edges, and shorten the average edge length. The latter is an indicator for the accuracy, i.e. the

spatial resolution which can be realized on the network. As we do not impose requirements regarding the computational performance towards the developed method, we do not consider the networks in a simplified version. Typically applied network simplification algorithms remove all nodes that are not intersections or dead-ends and further consolidates intersections such that intersection-nodes (degree larger 2) within a given radial distance from each other are merged into their respective centroids. Both these would significantly affect the realism of the network. Furthermore, one of the main attributes of interest are bike lanes and paths, which are often only present on certain segments (i.e. an edge in a series of edges) along a road without intersections.

network	layer	nodes	edges	mean edge length [m]	total length [km]
C	car	31,212	54,801	21.5	1,180
CB	car, bike	84,711	89,800	18.4	1,653
CBF	car, bike, foot	134,095	150,572	16.4	2,503

Table 1: Network (layer combination) characteristics.

The networks are further enriched with several empirically relevant choice attributes which are shown in Figure 6 in the Appendix. Figure 6a shows the existing cycling infrastructure (state 2017, network coverage 6.8%) in blue, including cycling lanes (total length 82.9km), defined as marked lanes on typical motorized street without physical barrier, and cycle paths (total length 88.3km), define as any path, track or lane physically separated from the motorized (not foot) infrastructure. The green segments represent the additionally envisioned cycling infrastructure presented under the current "Masterplan Velo 2030" (total additional length of approx. 120km) of the city of Zurich. The current and future cycling network are based on non-OSM networks (e.g., TomTom data) which do not geographically match to those used here. This makes it hard to automate the mapping to our network, and resulted in a manual annotation. Figure 6b shows the CBF network with a color coding depending on the absolute gradient. The corresponding elevation data is sourced from Google Maps Elevation API. The green color indicates rather flat edges (along the Limmat-Tal and around the Zurichsee), whereas the red color indicates steeper gradients due to the neighboring hills (Zueriberg, Uetliberg, Hoenggerberg).

Figure 6c shows the network parts for which traffic flow information is available. The data origins from the national traffic simulation model (ARE, 2017) and includes hourly counts per edge for passenger cars (LDV - light duty vehicles) as well as truck and delivery traffic with larger vehicles (HDV - heavy duty vehicles). The individual edges are mapped onto the car layer using distance-based map matching. After the mapping, the traffic count data has a network coverage of approx. 18.5%. Figure 6d shows the edges of the network for which speed limit information is available. The data is sourced from OSM and manually checked for plausibility

using reference data from the city of Zurich. Figure 6e and 6f show the point data for on-street parking spots (approx. 40,000 entries), as well as trees (and parks, approx. 80,000 entries). Both are publicly available and mapped to closest edge based on perpendicular distance using a 25m threshold value. After the mapping, on-street parking and tree information has a coverage of approx. 77.0% and 86.9% respectively.

3.3 Map matching

Both evaluated HMM implementations are based on the fundamental work of Newson and Krumm (2009) explained in Section 2. The HMM representation differs in both frameworks. Both use similar exponential functions for the transition- and emission-probabilities, and both allow to control for underlying noise in the data. The Leuven (LV) pipeline (Meert, 2020) assumes non-emitting states in the HMM graph. A non-emitting state (i.e. potential matching edge) is not associated with a matching observation, and hence provides a greater flexibility to handle outliers or inaccuracies in the GPS measurements. From each observation point, the edges within a radius of max_dist are considered as matching candidates, and their respective path probabilities on the HMM lattice is evaluated. The max_dist parameter has a large impact on computation times, as it defines how much of the graph is searched for each point that is to be matched. Considering the resolution of the network and the trajectories, the parameter was set to $max\_dist = 250m$. The GH framework (GraphHopper, 2020a) framework does not assume non-emitting states, every observation must hence be matched to a certain edge. Furthermore, the candidate search from each respective point is done using a Djisktra-based routine engine. This candidate generation is not restricted through a parameter like in the LV approach. The framework allows to modify the cost function used for routing, e.g. controlling turn-around behavior or time/distance preferences. For our use case, the default bike profile was used. Finally, and opposed to the LV approach, GH uses an internal downsampling of the input trajectory. The trajectories are typically filtered in such way, that the resulting distribution of the distance between two consecutive points is equal to at least double the assumed noise distribution.

The performance of map matching algorithms can be measured using different quantitative metrics. The all-over matching rate represents the share of successfully matched trajectories. For all of these, the accuracy can be measured using binary classification measures like the f-score. The underlying confusion matrix is derived using a distance threshold which determines if any GPS point is correctly classified, i.e. matched. Another accuracy indicator is the divergences for each trajectory segment, defined as average distance between each GPS point and its matched network edge. Finally, we also evaluate how the mapped length- and speed-distributions differ from the raw data. The computational performance of the algorithms is not considered.

3.3.1 Evaluation

The main parameter which is evaluated is the standard deviation of the assumed measurement noise, which can be controlled for in both frameworks. All other parameters, i.e. the transition- and emission-probability functions, or the non-emitting state related parameters in the LV pipeline are kept to default. Table 2 shows the performance metrics for the different network and algorithm configurations evaluated on a sub-sample of 1.000 trajectories. It can generally be said, that both pipelines have comparable performance and behave similarly w.r.t to the type of network and noise assumption. The performance in both pipelines improves with the network size, especially regarding divergence, f-score and speed-delta. In addition to the car layer, the foot layer adds more performance w.r.t. to the quantitative metrics than the bicycle one. This indicates that a substantial amount of the pedestrian (OSM-tag "footway" and "pedestrian") network is used in the observed trajectories. The noise parameter was tested in a range of 10m to 150m. For both pipelines the best results were achieved with a value of 50m. Too small values result in more failed and less accurate matches. Too large values typically result in more matches however with clearly increased divergence, decreased f-score.

algorithm	network	std. noise [m]	match rate [%]	mean divergence [m]	mean f-score	mean speed delta [km/h]
LV	C	50	61.2	40.5	0.67	-4.2
LV	CB	50	66.3	38.7	0.81	-2.8
LV	CBF	50	91.2	24.9	0.87	-0.7
GH	C	50	94.3	63.2	0.71	7.9
GH	CB	50	100	54.2	0.77	3.6
GH	CBF	50	100	40.7	0.80	-0.16
LV*	CBF	50	86.3	20.3	0.89	-0.4
GH*	CBF	50	86.6	34.3	0.86	0.7

Table 2: Quantitative results of different network and algorithm combinations.

It must be noted that the GH pipeline almost always finds a solution through the more sophisticated and unconstrained candidate generation. However, considering the trip-level accuracy metrics reveals that many results from the GH pipeline have bad and/or implausible results. An example is depicted in Figure 4 in the Appendix, for which the average matched speed is more than tripled (22 vs. 77km/h), and for which the LV algorithm could not find a solution. The results of both algorithms are hence filtered post matching based on the three defined metrics from above. Matching results with divergence values of 100m and higher or f-score values under 0.4, or absolute speed deltas of more than 5km/h are neglected and considered unsuccessful matches. The resulting performances are indicated on the bottom of Table 2, annotated by ∗. It can be

seen that the LV pipeline generates slightly better results, especially regarding divergence and speed delta. The GH pipeline tends to generate too fast, i.e. too long, matches, while speed delta from LV are negative and smaller.

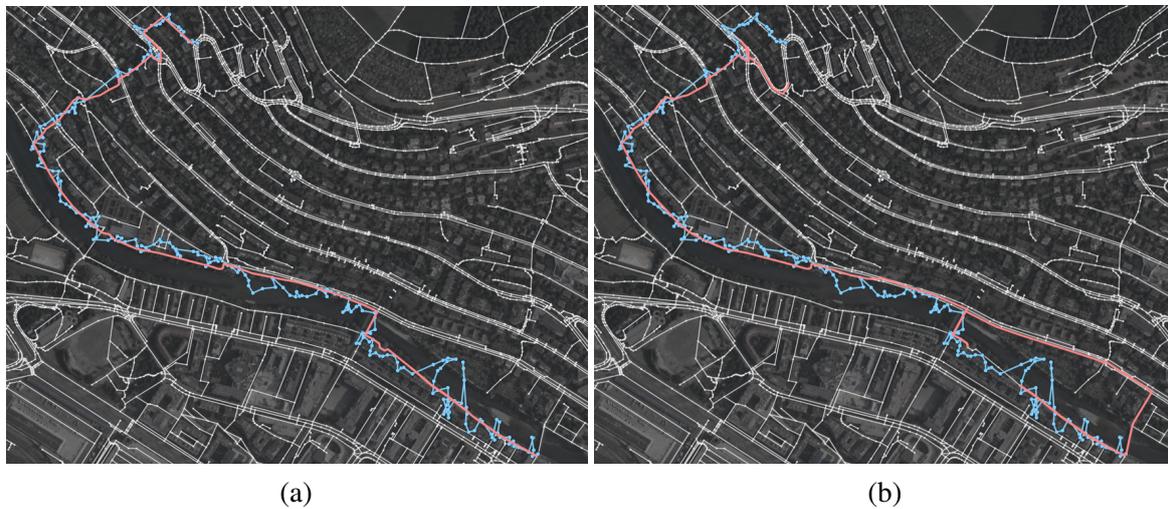| (a) | (b) |

Figure 2: Matching results from the LV (a, left) and GH pipeline (b, right).

After the filtering, the best combinations of hyper-parameters and network from both pipelines are further evaluated on a qualitative basis. Notable are the differences in routing behavior, like shown in Figure 2a and 2b (as well as Figure 5 in the Appendix). The right shows results from GH, the left from the LV pipeline. It can be seen that the paths matched by the GH pipeline take detours to cover (or "reach") (what seem to be) inaccurate GPS points, while deviating from the intuitive direction. This can be explained by the routine engine used by GH, which uses a bike-specific turn-around penalties (which are typically lower than for cars), as well as the inferior flexibility of the HMM representation which does not use non-emitting states and has to match every observation to an edge in the graph. This pattern can be observed in multiple results, and also explains the rather large positive speed delta values in Table 2.

3.3.2 Enrichment

The matching results allow it to enrich each trajectory by extracting the corresponding network attributes for the detected paths. The enrichment which results from the algorithm, parameter and network configuration are shown for the best combinations from the LV pipeline. The comparison further includes the best LV combination with a modified enrichment method, as well as the best GH combination as reference. The evaluated attributes include those which do not have a 100% coverage in the enriched network, i.e. traffic, trees and parking, speed limit as well as current and future cycling infrastructure. The combinations of algorithm and

network discussed in the previous section were evaluated in order to maximize the accuracy. The qualitative evaluation however showed, that the spatial resolution of the GPS measurements does for instance not allow to detect the side of the road (two lanes with same direction) or a sidewalk over a parallel running road. The latter case occurs often in especially dense parts of the network with all three car-, bike- and pedestrian-layer present. Figure 4 shows an example, i.e. the matched path (pink) of a GPS-series (blue, west to east) along with the underlying network (white/green). The three parallel lines running west to east represent two car lanes with opposing directions and a pedestrian lane on the southern end. It can be seen, that the matched path jumps between the lanes.
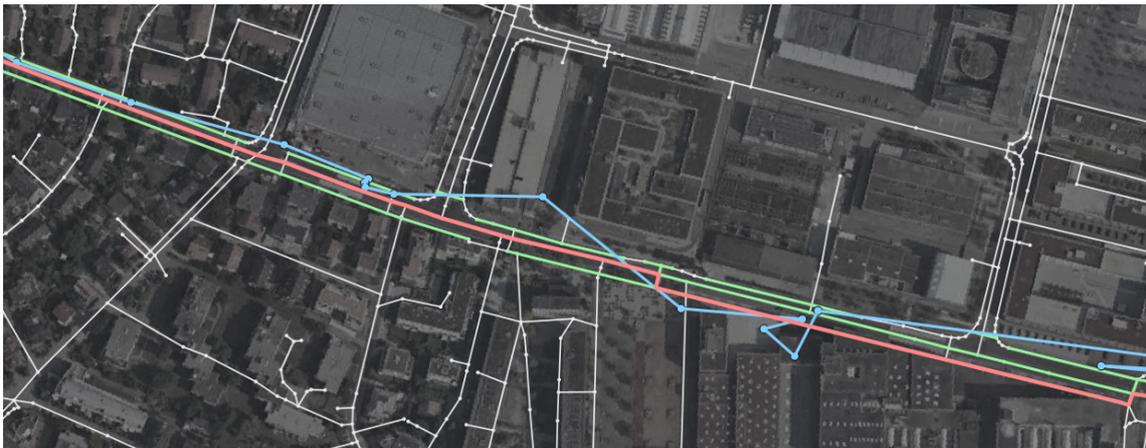

Figure 3: Matching results (pink) from the LV pipeline with corresponding raw GPS data (blue) and network (grey and green).

The rather low spatial resolution of the GPS measurements when compared to the network obviously leads to incorrect results and needs to be accounted for. However, even assuming noise-free, i.e. "perfect" GPS measurements, considering the matched edge independently from the surrounding unmatched ones would lead to a wrong and/or incomplete representation of the choice situation. From a choice perspective, it is for instance crucial to know if a bike lane runs in a pedestrian zone or next to a road with high traffic. To account for both these problems, a refined enrichment method is developed which takes into account parallel running edges for any given matched edge. Said method creates a sub-network with all edges in the surrounding of a given radius for each respective matched edge (perpendicular distance threshold of 30 meters). These edges are then filtered out if their direction (or their opposite direction, measured as azimuth) lies within a +/-2.5 deg. range of the one from the matched edge. The results are shown in Figure 4, where the green edges are those additionally attributed to the originally matched path (pink). A quantitative comparison of this modified and the regular enriched results is given in Table 3.

The degree to which the matched paths can be considered as enriched, is quantified using the mean share of path distance which is attributed with each respective attribute. The resulting

attribute/enrichment	network coverage [%]	CBF (LV) [%]	CBF (LV, ref) [%]	CBF (GH) [%]
LDV/h	18.4	48.8	60.1	28.5
HDV/h	17.8	50.6	59.4	31.9
trees/km	86.9	29.4	41.3	27.4
parking/km	77.0	21.1	25.3	22.5
speed limit	28.2	61.9	70.6	52.0
cur. lanes	3.3	13.7	16.5	11.1
cur. paths	3.5	10.3	12.9	8.2
fut. lanes/paths	4.5	16.2	19.7	14.2

Table 3: Quantitative evaluation of the attribute coverage of matched paths.

distribution of OSM road types for the four combination as well as the general CBF network is shown in Figure 10 in the Appendix. The results in Table 3 indicate that the GH pipeline generates a significantly less "rich" dataset, with some indicators about half as large as the CBF (LV) combination. This is due to the comparably high share of pedestrian edges in the GH matching results. Accordingly, the difference between LV and GH is specfcially pronounced for attributes which are mapped to the car layer (traffic, speed limits) and less visible for those mapped to the pedestrian layer (trees, parking). The refined enrichment method which accounts for surrounding parallel lanes drastically increases all enrichment indicators (approx. 40% increase for trees), and increases the captured share of the pedestrian layer in the LV matching results as indicated in Figure 10. The effect on the cycling infrastructure (approx. 20% increase in coverage) attributes is smaller due to the generally low network coverage.

4 Results

The evaluation in the previous chapters has been conducted with a subset of 1,000 trajectories for computational reasons. The following chapter presents the descriptive results of the complete 4,978 matched trajectories. The attributes considered include spatio-temporal characteristics (duration, length, speed) and those which are derived from the map matching-based enrichment. Table 5 and 6 in the Appendix additionally lists trip- and traveler-characteristics, those are however not further discussed here. We consider all cycling modes holistically because of the too few reported e-bike trips and do not further decompose the results. The results of the CBF (LV) combination with the regular edge-based and the refined enrichment method are shown in Table 4. Note that the refined enrichment method is only applied to the non-topological attributes which do not have a complete network coverage. Where applicable, the attributes are compared to the raw pre-matching GPS data or values from other studies, specifically work from Menghini et al. (2010). Considering other studies is difficult due to the different geographical and local transport-system characteristics. The following discusses each attribute individually. The normalized density of a Gaussian-kernel transformation for each respective distribution is shown in Figure 7, 8 and 9 in the Appendix. The trip duration is kept constant throughout the whole data processing.

The distribution of the original and the matched trip length is shown in Figure 7a. After cleaning in-plausible matching results, both distribution match w.r.t. mean and standard deviation and are right-skewed. The trip length is longer when compared to results from Menghini et al. (2010) which had a mean length of only 1km. The MZMV indicates cycling trip lengths of around 3.2km for bikes and around 5km for e-bikes. The mean speeds in the raw data are slightly faster than those found in Menghini et al. (2010) (only approx. 11km/h). The MZMV reports mean speeds of 13.3km/h for bikes and 17km/h for e-bikes, value ranges which have recently been reported by Dane et al. (2019). Both the raw and matched speed distributions (see Figure 7c) are comparable w.r.t. mean and standard deviation and both distributions have a rather symmetrical shape. The mean speed values suggest that a certain amount of e-bike trips is present in the raw data. The mode detection of the CMD app does not include e-bikes. The (partly user-corrected) labelling indicates a share of e-bikes of only 0.2% which is unrealistic considering the just evaluated speed characteristics, as well as the ownership indicators from the MOBIS-COVID sample and MZMV. The mean gradient of the matched paths is approx. 3%, the mean maximum positive gradient approx. 18.3% (see Figure 7e and 7f). The work from Menghini et al. (2010) which covered the same geographical city boundaries indicates significantly lower values (approx. 1% mean grade and 3% maximum grade). While the distribution of the avg. gradient has a symmetric, unimodal shape, the one of the mean max. positive gradient is right-skewed and has a rather bimodal shape. The latter suggests the presence of two distinct groups, a further indicator

network/algorithm	raw GPS-data	CBF (LV)	CBF(LV, ref)
mean duration [min]	9.2 (0)	-	-
mean length [km]	2.3 (1.5)	2.2 (1.5)	-
mean speed [km/h]	16.8 (6.5)	16.1 (6.8)	-
mean routing speed [km/h]		15.1 (1.9)	-
mean gradient [%]		0.3 (1.1)	-
mean max. pos. gradient [%]		18.3 (15.7)	-
mean traffic [LDV/h]		2,234.3 (1,716.8)	1,944.1 (1,526.7)
mean traffic [HDV/h]		238.6 (218.7)	208.7 (194.9)
mean trees [n/km]		32.8 (28.5)	34.8 (26.1)
mean on-street parking [n/km]		32.3 (31.8)	27.7 (25.7)
mean max. speed limit [km/h]		48.2 (7.2)	49.3 (5.7)
mean cur. lanes [%]		13.7 (16.8)	16.5 (19.1)
mean cur. paths [%]		10.3 (16.0)	12.9 (17.8)
mean fut. lanes/paths [%]		16.2 (17.4)	19.7 (18.9)

Table 4: Descriptive results from the regular and the refined enrichment method.

for a substantial amount of electrified cycling trips in the data. The traffic, trees and on-street parking indicators do not show considerable differences between the two applied enrichment methods. The distributions for all four attributes have similar right-skewed shapes for both methods, with the refined method producing slightly narrower distributions. The max. speed limit density (Figure 9b) shows a major peak at 50km/h, which is the general speed limit for the city of Zurich. Even-though not clearly visible but expected, the regular enrichment shows peaks at 0/30km/h, while the refined enrichment also shows peaks at 60/80km/h. The mean

share of current and future cycling infrastructure used in the observed trajectories is around 15% and 18% respectively. This is considerable smaller than values of almost 75% reported in Menghini et al. (2010). Analog to the previously discussed differences, this can be related to the different goals of the respective data collection method which generated a rather narrow and biased sample as opposed to the MOBIS-COVID data. The lanes are slightly more used than the paths even-though they have a lower network coverage, the differences are however rather small. The current and future infrastructure have similar marginal rates of matched, i.e. used share of cycling infrastructure in relation to the network coverage, and can hence be considered to be comparably in their effectiveness. The indicators for current and future cycling infrastructure both show similar right-skewed distributions when comparing the regular and refined enrichment method. As opposed to the previous attributes, here the refined method produces slightly wider and more symmetrical distributions.

5 Conclusion

This work presented the preliminary results of a descriptive route choice analysis of cycling trips in Zurich. The trajectory data was sourced from the MOBIS-COVID dataset and matched to an OSM network using an HMM-based map-matching pipeline. Said pipeline was optimized regarding its hyper-parameters as well as the underlying network. The results show, that both evaluated implementations of the HMM-based map matching behave similar to changes in parameters of network configurations. The quantitatively measured matching performance of both pipelines intuitively increases with the network size. For our data, the qualitative evaluation however shows a superiority of the LV pipeline. This is assumed to be linked to the simpler candidate generation technique, which only consists of a random search of the HMM-graph, as opposed to more sophisticated routing using cost-functions in the GH framework. Following both the quantitative and qualitative evaluation, the results are further compared regarding their resulting enrichment. An important finding includes, that the accuracy of the smartphone-collected GPS measurement is not high enough to fully leverage the spatial resolution of the OSM network. Especially when multiple car- and pedestrian edges run parallel to each other, the results of the matching are rather arbitrary and an exact path identification is not possible. To overcome this problem, a refined enrichment method is proposed, which considers all parallel network edges within a specific radius. Not only does this increase the resulting coverage of specifically sparse network attributes in the matched paths, but it also provides a more complete representation of the choice situation along each path segment. Future work will focus on modelling the data using suited route choice models based on discrete choice frameworks. The results will help to evaluate the currently envisioned cycling infrastructure and hence help to create an urgently needed more sustainable transport system.
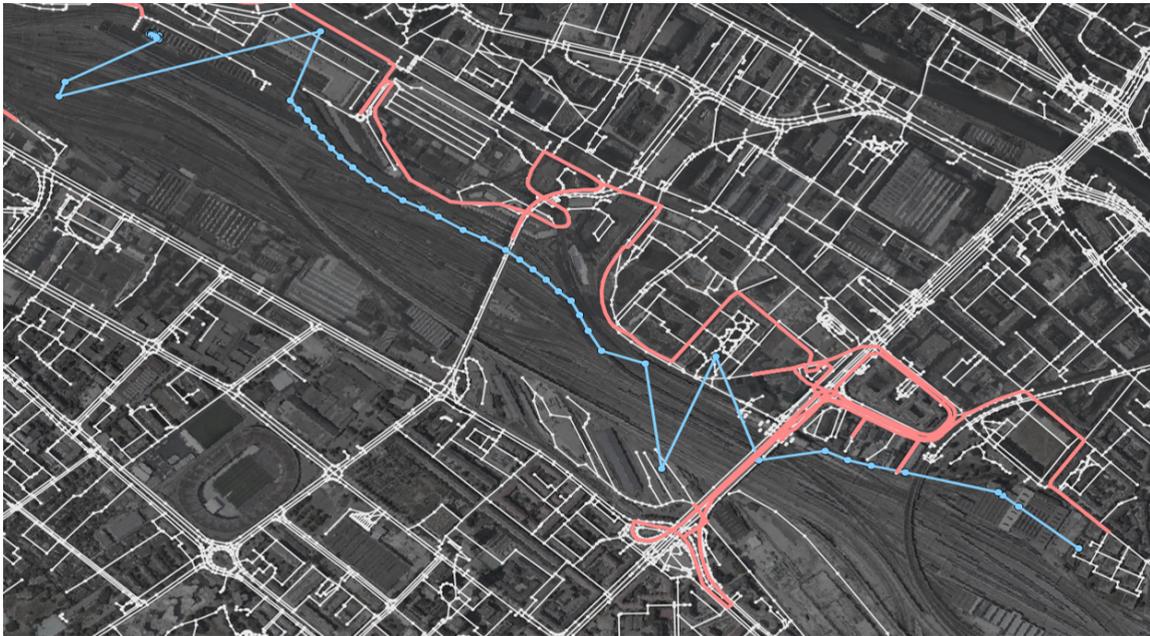
6 Appendix

Figure 4: In-plausible matching results from the GH pipeline. Raw GPS in blue, matched path in pink.

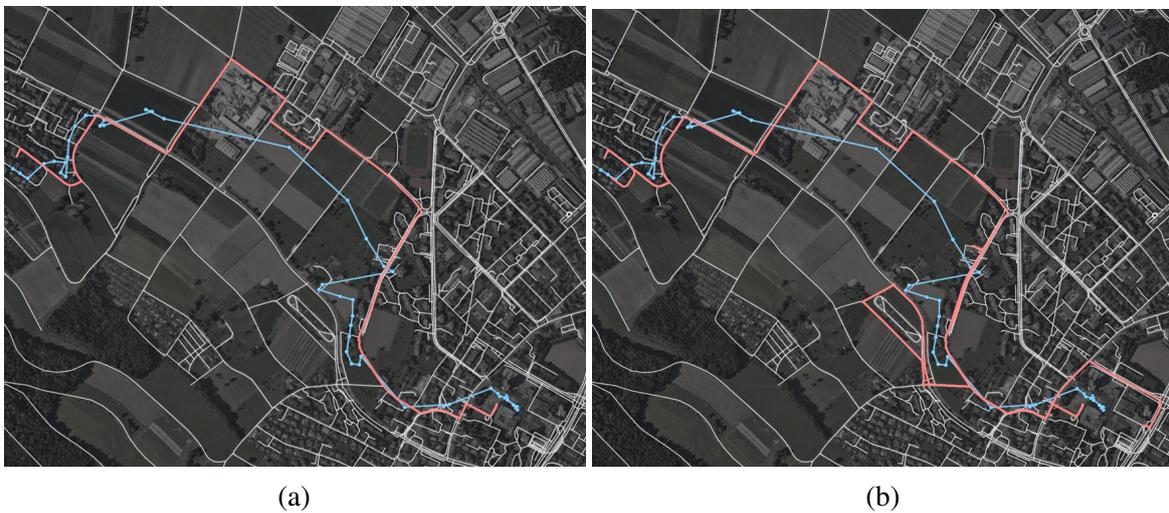(a)   (b)

Figure 5: Matching results from the LV (a, left) and the GH pipeline (b, right). Raw GPS in blue, matched path in pink.
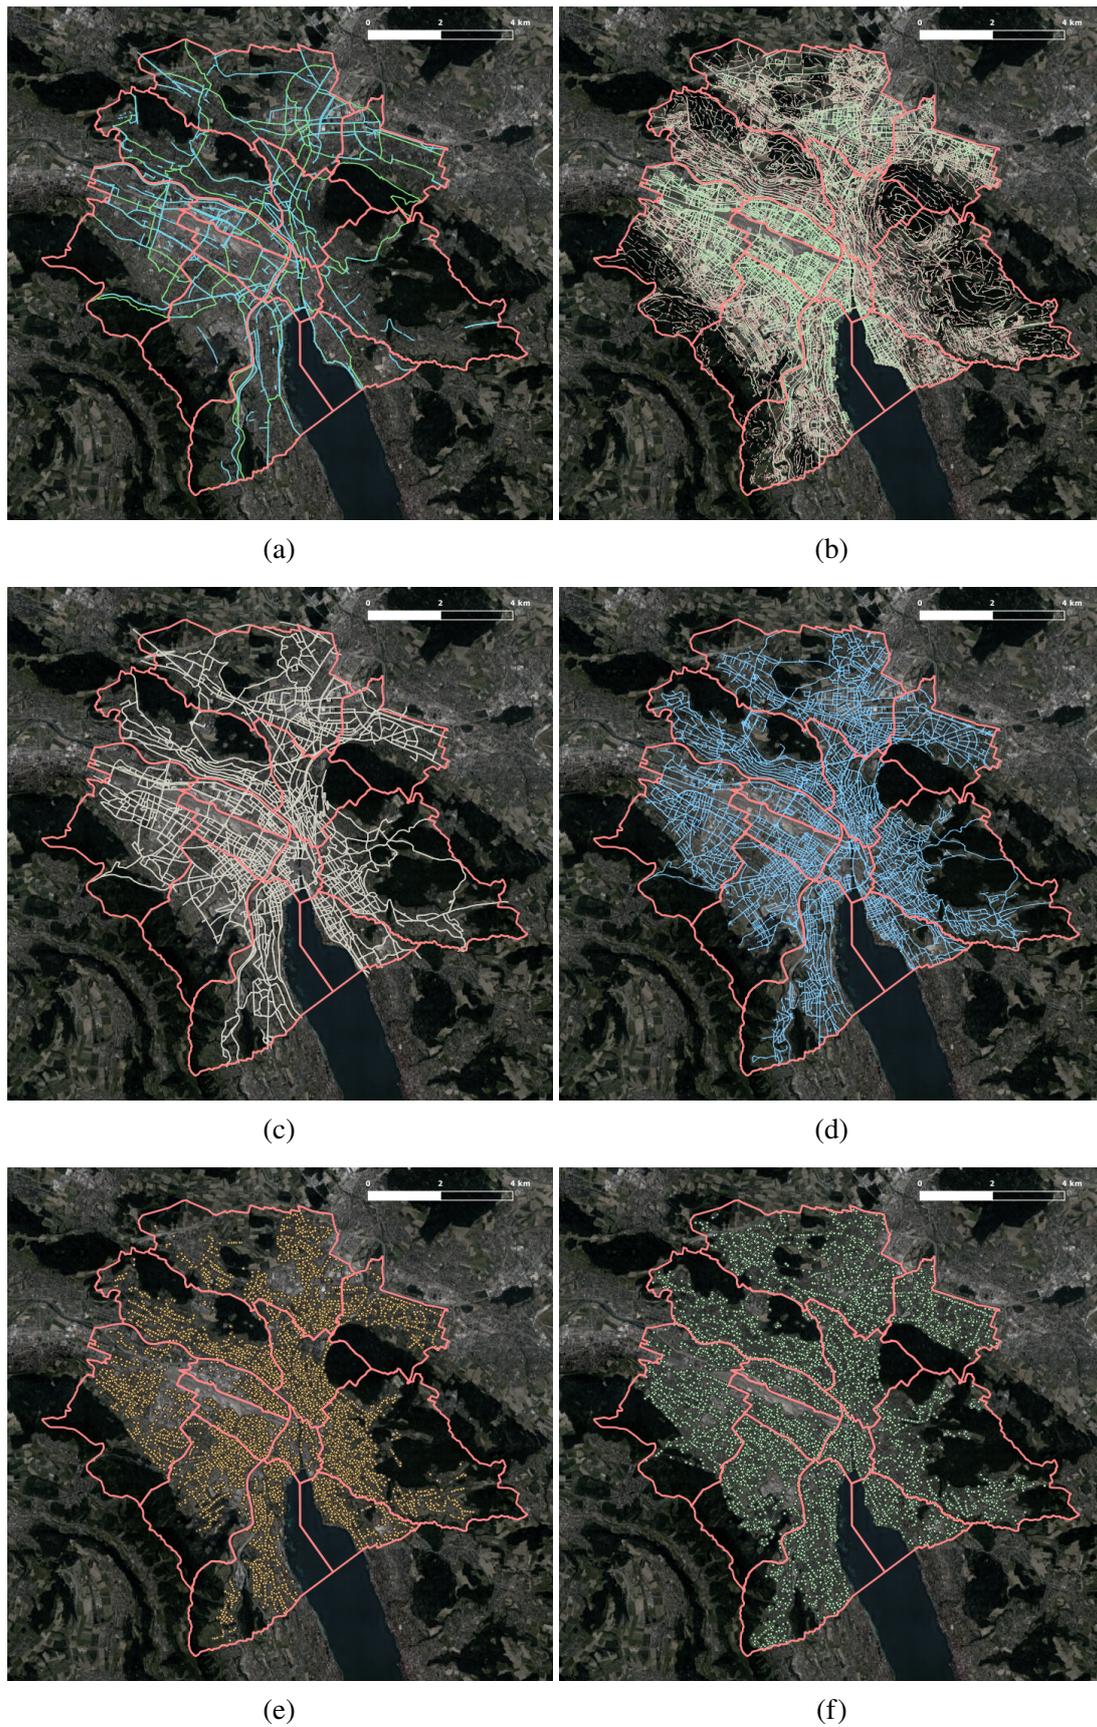
Figure 6: Study area (pink) with cycling infrastructure (a, blue and green), gradients (b, green/flat to red/steep), traffic count data (c, grey), speed-limit (d, blue), on-street parking (e, orange) and trees (f, green).
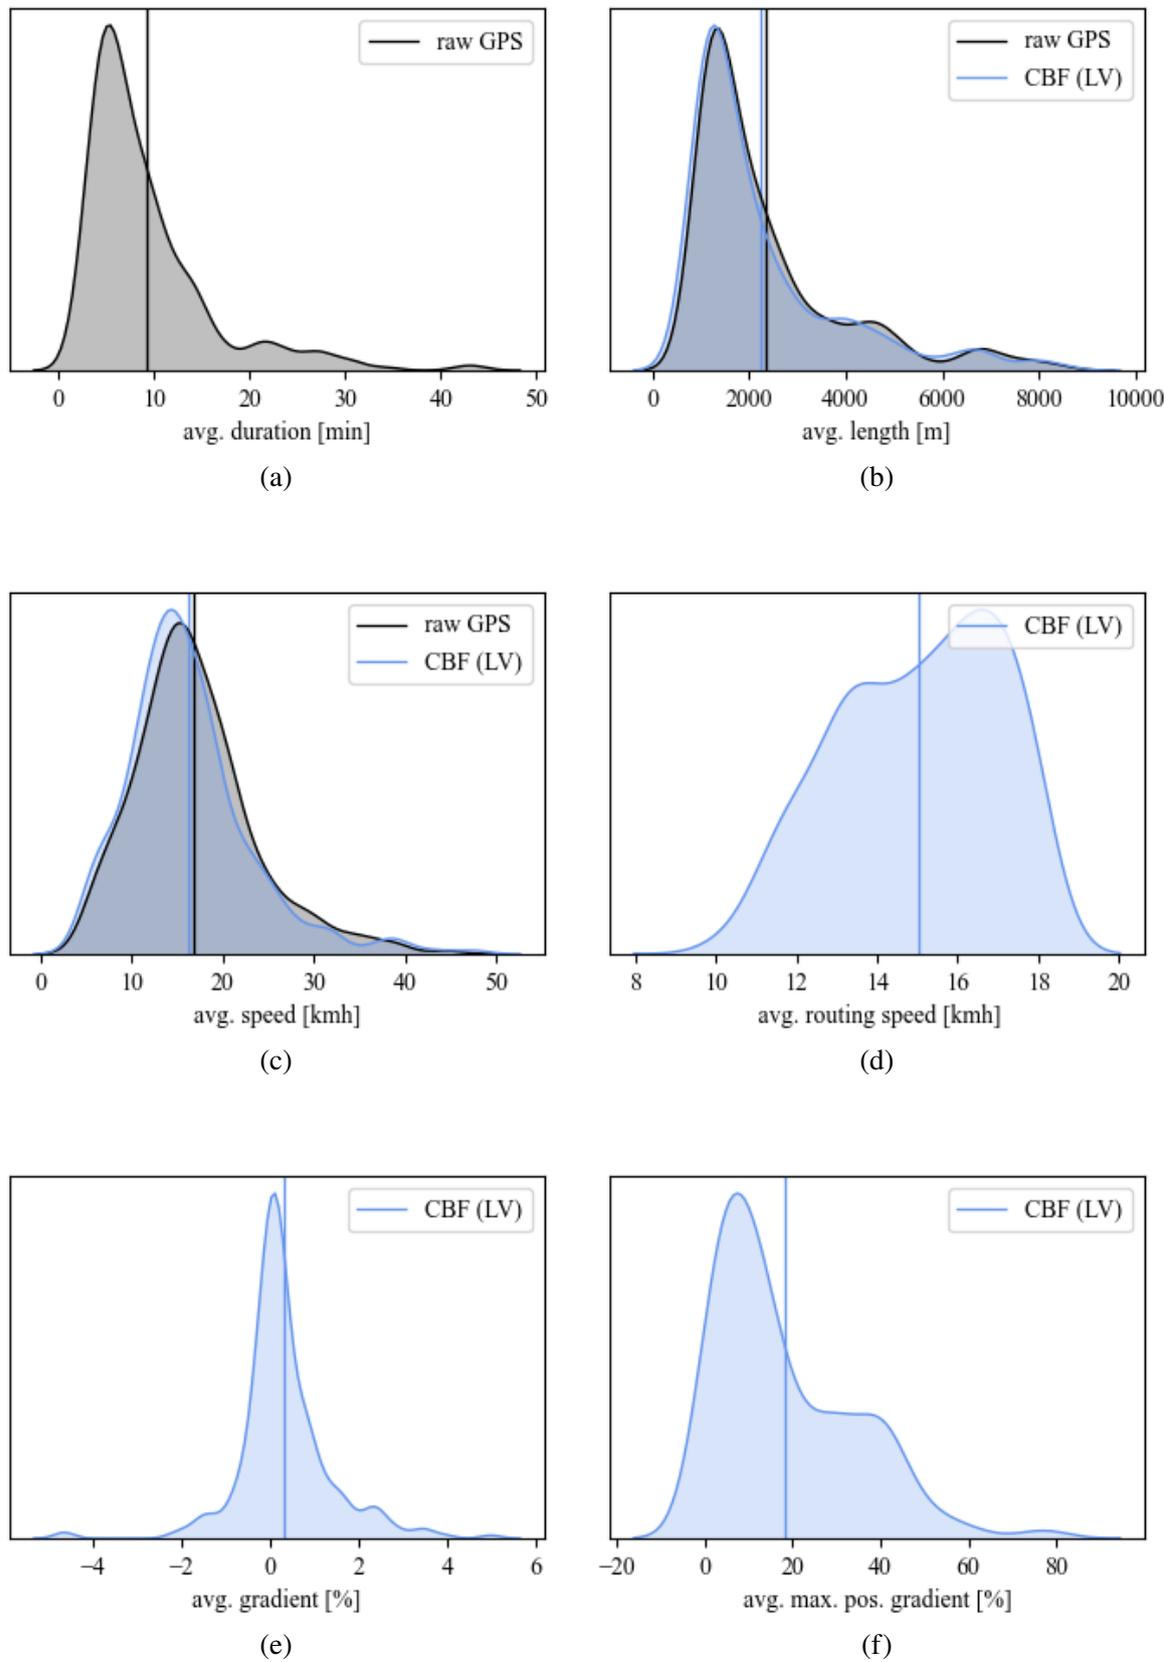
Figure 7: Normalized densities of duration (a), length (b), speed (c), routing speed (d), gradient (e) and max. pos. gradient (f).
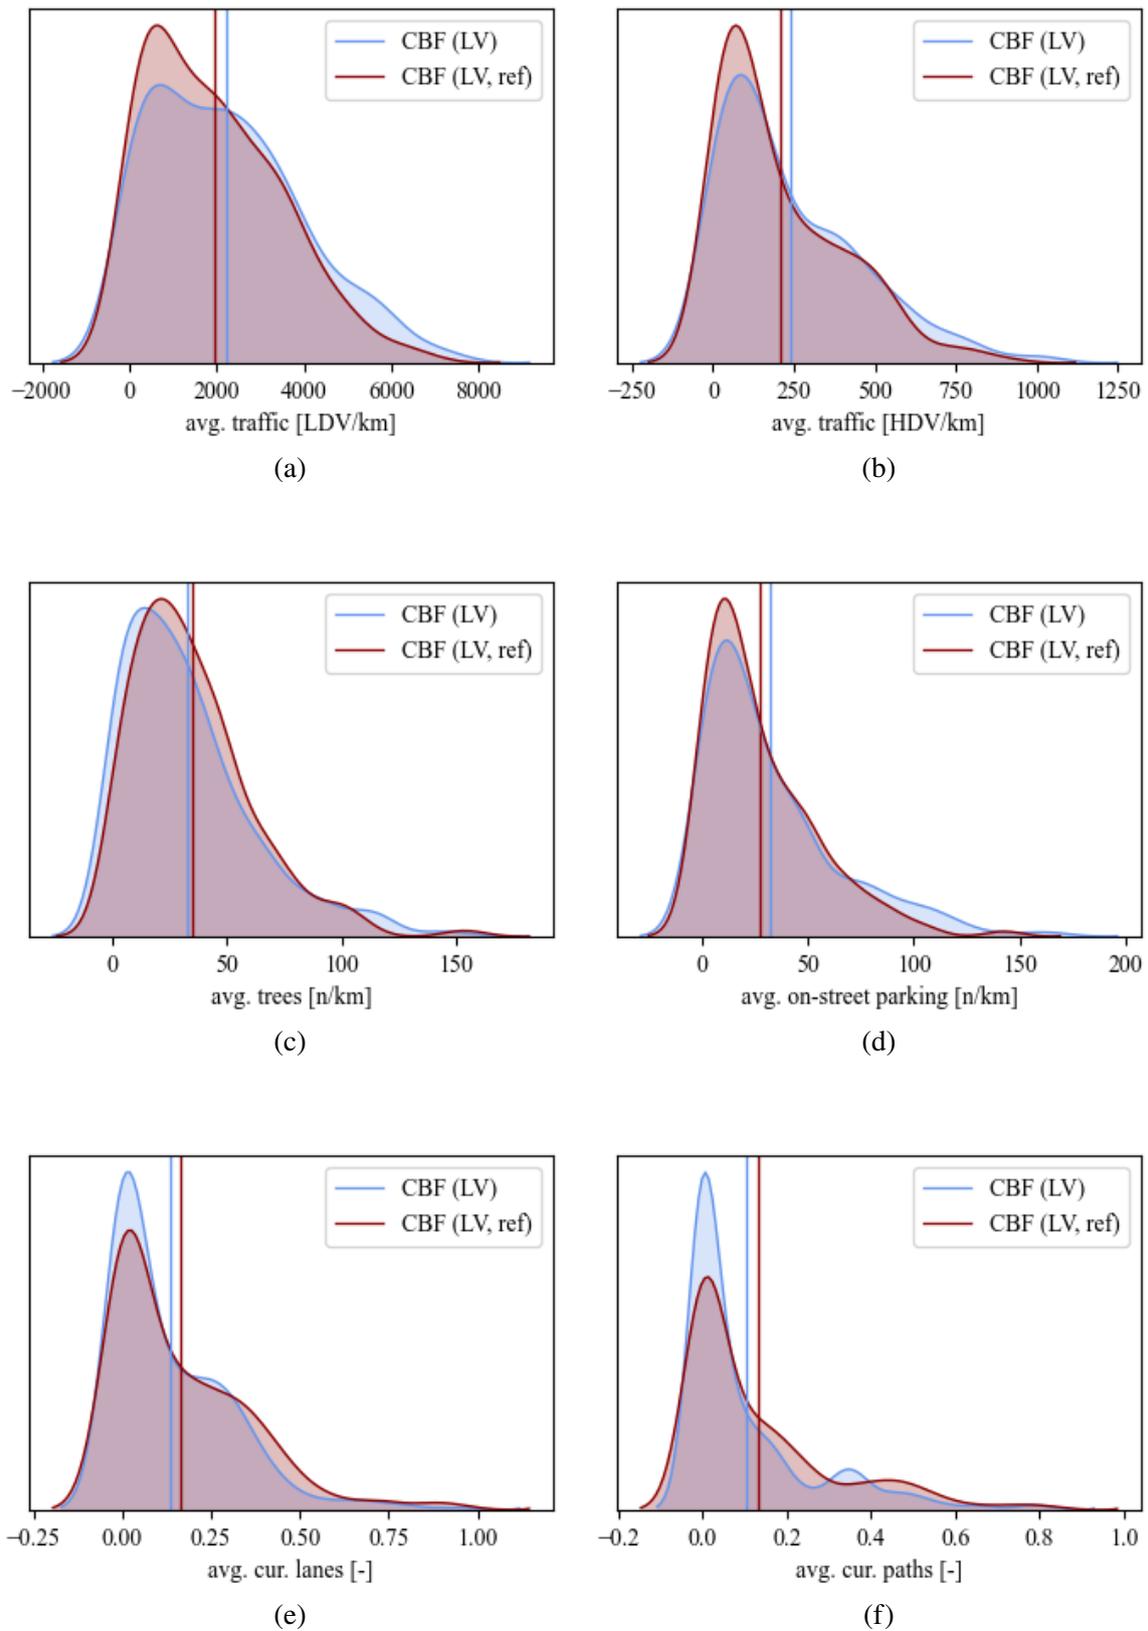
Figure 8: Normalized densities of LDV traffic (a), HDV traffic (b), trees (c), on-street parking (d) as well as current lanes (e) and current paths (f).
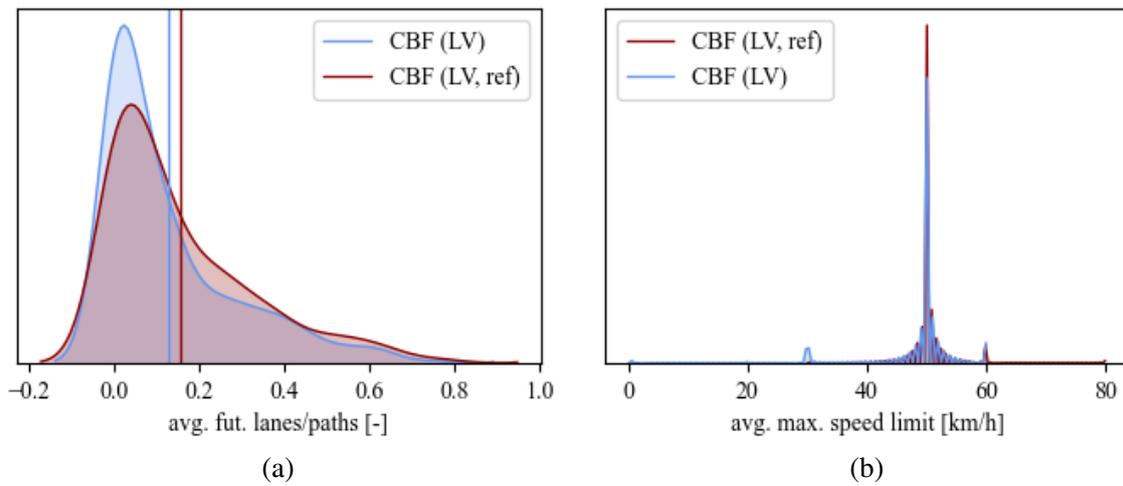
Figure 9: Normalized densities of future lanes/paths (a) and max. speed limit (b).
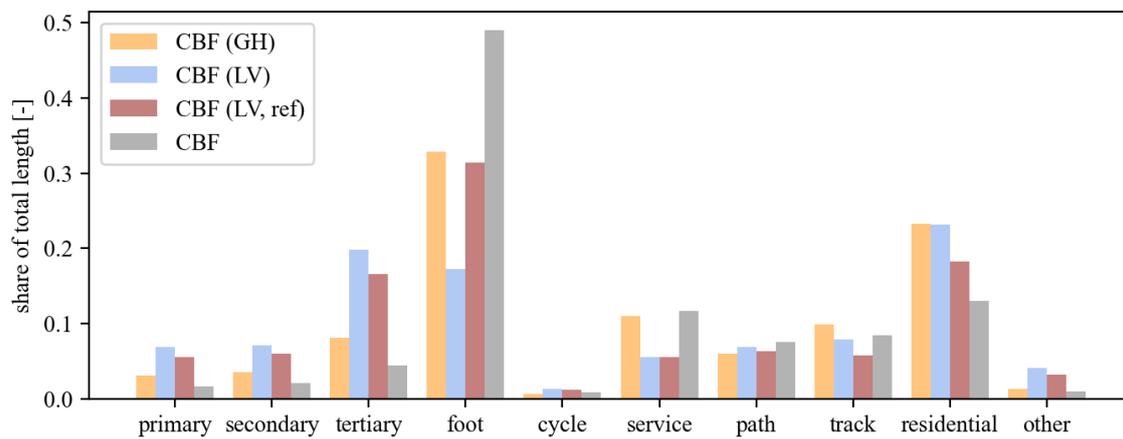

Figure 10: Distribution of OSM road types from different evaluated configurations.

variable		raw GPS-data [%] n=5,656	CBF (LV) [%] n=4,978
trip purpose	home	31.5	37.7
	other	16.6	11.7
	work	15.8	16.8
	shopping	10.5	13.0
	assistance	1.7	1.6
	education	1.6	2.0
	errand	0.6	1.0
	leisure	21.7	16.8

Table 5: Distribution of trip purposes from the raw and matched data.

variable			census [%]	raw GPS-data [%] n=367	CBF (LV) [%] n=359
traveler	age	under 18	13.2	-	-
		18-25	9.0	17.2	15.4
		25-45	29.6	40.1	43.5
		45-65	29.6	42.7	41.1
		over 65	18.5	-	-
	gender	male	49.3	51.0	51.2
		female	50.7	49.0	48.8
	income	not said	20.7	8.6	13.5
		low	17.8	4.9	6.0
		medium	50.2	61.1	57.4
		high	11.3	25.4	23.1
	education	mandatory	19.3	6.3	6.2
		secondary	49.5	45.3	48.6
		higher	31.2	48.4	45.2

Table 6: Distribution of socio-demographic attributes from the respondents with MZMV data as reference.

7 References

ARE (2017) *Nationales Personenverkehrsmodell des Bundes*, ARE - Bundesamt fuer Raumentwicklung.

Axhausen, K. W., A. Horni and K. Nagel (2016) *The multi-agent transport simulation MATSim*, Ubiquity Press.

BFS (2017) *Verkehrsverhalten der Bevoelkerung*, BFS - Bundesamt fuer Statistik.

Blondel, B., C. Mispelon and J. Ferguson (2011) Cycle more often 2 cool down the planet, *Brussels: European Cyclists Federation*.

Broach, J., J. Dill and J. Gliebe (2012) Where do cyclists ride? A route choice model developed with revealed preference gps data, *Transportation Research Part A: Policy and Practice*, **46** (10) 1730 – 1740.

Casello, J. M. and V. Usyukov (2014) Modeling cyclists' route choice based on gps data, *Transportation Research Record*, **2430** (1) 155–161.

Cazzola, P. and P. Crist (2020) Good to go? Assessing the environmental performance of new mobility.

Chao, P., Y. Xu, W. Hua and X. Zhou (2020) A survey on map-matching algorithms, paper presented at the *Australasian Database Conference*, 121–133.

Dane, G., T. Feng, F. Luub and T. Arentze (2019) Route choice decisions of e-bike users: Analysis of gps tracking data in the Netherlands, paper presented at the *The Annual International Conference on Geographic Information Science*, 109–124.

de Freitas, L. M., H. Becker, M. Zimmermann and K. W. Axhausen (2019) Modelling intermodal travel in Switzerland: A recursive logit approach, *Transportation Research Part A: Policy and Practice*, **119**, 200–213.

Forney, G. (1973) The viterbi algorithm, *Proceedings of the IEEE*, **61** (3) 268–278.

GraphHopper (2020a) Graphhopper: Map-matching library, `https://github.com/graphhopper/map-matching`.

GraphHopper (2020b) Graphhopper: Open source routing software, `https://github.com/graphhopper/graphhopper`.

Hood, J., E. Sall and B. Charlton (2011) A gps-based bicycle route choice model for San Francisco, California, *Transportation letters*, **3** (1) 63–75.

Jensen, C. S. and N. Tradišauskas (2009) *Map Matching*, 1692–1696, Springer US.

Khatri, R., C. R. Cherry, S. S. Nambisan and L. D. Han (2016) Modeling route choice of utilitarian bikeshare users with gps data, *Transportation research record*, **2587** (1) 141–149.

Lu, W., D. M. Scott and R. Dalumpines (2018) Understanding bike share cyclist route choice using gps data: Comparing dominant routes and shortest paths, *Journal of transport geography*, **71**, 172–181.

Meert, W. (2020) Leuvenmapmatching, `https://github.com/wannesm/LeuvenMapMatching`.

Menghini, G., N. Carrasco, N. Schüssler and K. W. Axhausen (2010) Route choice of cyclists in zurich, *Transportation research part A: policy and practice*, **44** (9) 754–765.

Misra, A. and K. Watkins (2018) Modeling cyclist route choice using revealed preference data: an age and gender perspective, *Transportation Research Record*, **2672** (3) 145–154.

Molloy, J., T. Schatzmann, B. Schoeman, C. Tchervenkov, B. Hintermann and K. W. Axhausen (2021) Observed impacts of the covid-19 first wave on travel behaviour in Switzerland based on a large gps panel, *Transport Policy*, **104**, 43–51.

Newson, P. and J. Krumm (2009) Hidden markov map matching through noise and sparseness, paper presented at the *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*.

Schuessler, N. and K. W. Axhausen (2009) Processing raw data from global positioning systems without additional information, *Transportation Research Record*, **2105** (1) 28–36.

Sims, R., R. Schaeffer, F. Creutzig, X. Cruz-Nunez, M. Dagosto, D. Dimitriu and R. Sausen (2014) IPCC Fifth Assessment Report (AR5)-Chapter 8: Transport.

Zimmermann, M., T. Mai and E. Frejinger (2017) Bike route choice modeling using gps data without choice sets of paths, *Transportation research part C: emerging technologies*, **75**, 183–196.