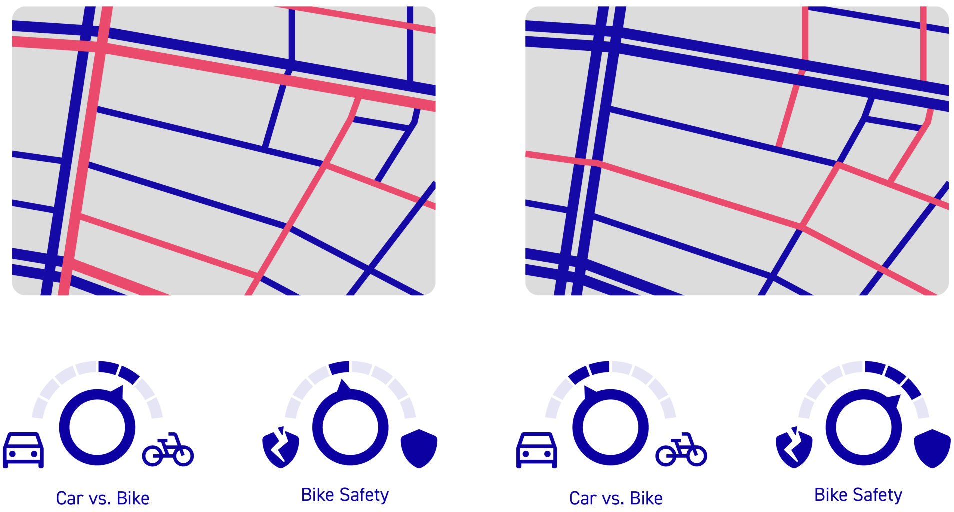

Bike network planning in limited urban space

Nina Wiedemann

Christian Nöbel

Henry Martin

Lukas Ballo

Martin Raubal

STRC conference paper 2024 **April 28, 2024**

STRC | 24th Swiss Transport Research Conference
Monte Verità / Ascona, May 15-17, 2024

Bike network planning in limited urban space

Nina Wiedemann
MIE Lab
ETH Zurich
nwiedemann@ethz.ch

Christian Nöbel
IFOR
ETH Zurich
christian.noebel@math.ethz.ch

Henry Martin
MIE Lab
ETH Zurich
martinhe@ethz.ch

Lukas Ballo
IVT
ETH Zurich
lukas.ballo@ivt.baug.ethz.ch

Martin Raubal
MIE Lab
ETH Zurich
mraubal@ethz.ch

April 28, 2024

Abstract

The lack of cycling infrastructure in urban environments hinders the adoption of cycling as a viable mode for commuting, despite the evident benefits of (e-)bikes as sustainable, efficient, and health-promoting transportation modes. Bike network planning is a tedious process, relying on heuristic computational methods that frequently overlook the broader implications of introducing new cycling infrastructure, in particular the necessity to repurpose car lanes. In this work, we call for optimizing the trade-off between bike and car networks, effectively pushing for Pareto optimality. This shift in perspective gives rise to a novel linear programming formulation towards optimal bike network allocation. Our experiments on six real urban street networks testify the effectiveness and superiority of this optimization approach compared to heuristic methods. In particular, the framework provides stakeholders with a range of lane reallocation scenarios, illustrating potential bike network enhancements and their implications for car infrastructure. Crucially, our approach is adaptable to various bikeability and car accessibility evaluation criteria, making our tool a highly flexible and scalable resource for urban planning. This paper presents an advanced decision-support framework that can significantly aid urban planners in making informed decisions on cycling infrastructure development.

Keywords

cycling infrastructure; urban planning; discrete optimization; linear programming

Suggested Citation

Wiedemann, Nina, Christian Nöbel, Henry Martin, Lukas Ballo, and Martin Raubal. "Bike network planning in limited urban space" *24th Swiss Transport Research Conference*

Acknowledgements

We thank Aurélien Borgeaud for his contribution in the context of a semester project. This project is part of the E-Bike City project funded by the Department of Civil and Environmental Engineering (D-BAUG) at ETH Zurich and the Swiss Federal Office of Energy (BFE).

Contents

Acknowledgements . . . 1
List of Tables . . . 2
List of Figures . . . 2
1 Introduction . . . 3
2 Problem definition . . . 4
3 Methods . . . 5
3.1 An evaluation framework based on Pareto optimality . . . 5
3.2 A linear programming approach towards solving the BNAP . . . 6
3.2.1 Introducing bike, car and shared flow . . . 6
3.2.2 Constraining the space for bike and car lanes . . . 7
3.3 Linear relaxation . . . 8
4 Experiments . . . 8
5 Discussion . . . 11
6 References . . . 12
A Overview of all variables . . . 15
B Post-processing algorithm . . . 15
C Preprocessing of real network data . . . 16
D Heuristic baselines . . . 16

List of Tables

1 List of variables and explanation 15
2 Overview of real-world instances. The runtime for solving the LP one time is given as “Runtime optimization”, in contrast to the runtime for computing the whole Pareto frontier (“Runtime Pareto”). 17

List of Figures

1 Network layout of the real instances used to test our algorithm 9
2 Pareto optimality of bike networks. Top: Algorithms are compared by their Pareto frontier. In five out of six instances, our linear programming approach outperforms methods based on the betweenness centrality. Bottom: Each point on the Pareto frontiers (top) corresponds to one plausible street network. Three examples in Cambridge MAare shown, where the bike networks differ dependent on the planning method and the number of allocated bike lanes. This is also reflected in the distance of shortest paths, where the existence of dedicated cycling infrastructure is rewarded in the perceived bike travel time. 10

1 Introduction

The transportation sector plays a pivotal role in combating climate change, accounting for around 20.7 % of CO_2 emissions worldwide (EDGAR/JRC, 2022). In urban environments, a viable alternative to motorized transport is cycling, promising sustainable traffic in addition to substantial health benefits (Oja *et al.*, 2011). The emergence of e-bikes in recent years has further democratized cycling for the general population. Nevertheless, in most cities cycling only accounts for a minority of transport activity so far. A crucial factor in the adoption of bicycle commuting is the availability and density of bike networks (Schoner and Levinson, 2014). Previous research has provided compelling evidence for the positive impacts of bike lane infrastructure on cycling (Buehler and Dill, 2016), encompassing its effects on public health (Mueller *et al.*, 2018), the importance of physical separation from car lanes (Fraser and Lock, 2011), and the incorporation of green spaces (Ta *et al.*, 2021). These findings collectively underscore the importance of well-designed bike infrastructure in promoting sustainable and healthy urban transportation choices.

Many cities around the world have recently promoted the construction of large-scale cycling infrastructure. Implementation often suffers from practical difficulties and a complex planning process that involves many stakeholders. Computational methods can serve as decision support systems, including topology-based methods aiming to improve the connectivity of bike networks (Natera Orozco *et al.*, 2020), cost-benefit analyses (Paulsen and Rich, 2023; Szell *et al.*, 2022), or data-driven planning methods based on data from bike sharing systems (Steinacker *et al.*, 2022; Duthie and Unnikrishnan, 2014; Bao *et al.*, 2017; Liu *et al.*, 2022), travel surveys (Mauttone *et al.*, 2017) or mobile phones (Olmos *et al.*, 2020). Methods range from simple heuristics, e.g. based on the betweenness centrality of edges in the network (Steinacker *et al.*, 2022; Ballo *et al.*, 2023), to linear programming (LP) (Duthie and Unnikrishnan, 2014; Lin and Yu, 2013) or mixed integer linear programming (MILP) (Liu *et al.*, 2022) approaches.

However, most approaches that propose large scale changes of the existing infrastructure do not consider the impact of these changes on other transport modes (Gerike *et al.*, 2022). As cities have a (mostly) fixed street infrastructure, an improved cycling infrastructure is only attainable by allocating existing streets or lanes from car to bike usage. Thus, every improvement in bike infrastructure comes at the cost of worsening the car network (Burke and Scott, 2016). As such an essential factor for the practical feasibility and public acceptance of a radical network restructuring in favour of cycling is its impact on the car network.

The contribution of this work is two-fold. First we propose a framework to evaluate bike network planning approaches that takes into account both the improvement of the bike infrastructure as well as the car reachability in the modified network. This is achieved by comparing planning methods in terms of their *Pareto frontier*, trading off the bikeability against the car travel times. We supplement this quality measure with a novel optimization framework for bike network planning. The framework can provide stakeholders with a multitude of possible reallocations, showing which improvements of the bike network can be achieved at what cost for the car infrastructure.

2 Problem definition

We propose a multi-modal view on bike network planning that considers the impact of new bike lanes on the car network. It is based on the following assumptions: 1) The input is a given street network of an urban area. It is not possible to build entirely new infrastructure, but the type and division of existing road space can be changed, including lane directions. 2) The allocation of bike lanes inevitably involves the reduction of street space available to other modes.

This planning problem, here termed the "bike network allocation problem" (BNAP), can be modeled as a graph division problem. The initial graph is a simplified version of the existing street network, corresponding to an undirected graph $G = (V, E)$, where the nodes V are intersections and the edges E are streets between two intersections. Geographic properties of the network are expressed in attributes of the street edges e, specifically the length of the street $d(e)$, its gradient $\delta(e)$, the speed limit $\theta(e)$ and its capacity $\lambda(e)$. The capacity can be set to the width of the street or the number of lanes.

Solving the BNAP involves dividing G into two graphs: the bike lane network $G_b = (V, E_b)$ and the car lane network $G_c = (V, E_c)$. Both are directed multi-graphs, since their edges now represent *lanes*. The car network G_c must be strongly connected; i.e., every node must be reachable from any other node, since disconnected subgraphs are unrealistic in an urban environment. The design of G_b and G_c is mainly constrained by the street capacities, ensuring that the car and bike lanes fit into the existing infrastructure:

$$\forall e = (u, v) \in E : \lambda^c_{(u,v)} + \lambda^c_{(v,u)} + 0.5(\lambda^b_{(u,v)} + \lambda^b_{(v,u)}) \leq \lambda(e) \tag{1}$$

Here, $\lambda^c_{(u,v)}$ denotes the capacity of a directed edge $e = (u, v)$ in the car network G_c and $\lambda^b_{(u,v)}$ the capacity in the bike network. If the edge (u, v) is not part of the car network, we define $\lambda^c_{u,v} = 0$ and accordingly for the bike network. Equation 1 defines a division of the total capacity of the undirected street into directed car and bike lanes. The bike capacities are multiplied by 0.5 to express the lower space necessary for bikes. We follow common guidelines for bicycle infrastructure that recommend a bike lane width of 1.5 m (Parkin, 2018; Yan *et al.*, 2018), corresponding to about half of a car lane.

3 Methods

3.1 An evaluation framework based on Pareto optimality

While a plethora of metrics has been proposed for evaluating "bikeability" (Weikl and Mayer, 2023; Grisiute *et al.*, 2023), they largely ignore the impact of proper bike lanes on motorized travel, with the exception of Burke and Scott (2016) who propose the Network Robustness Index to measure the effect of wider bike lanes on traffic. For simultaneously evaluating the goodness of the bike and car network, we leverage the concept of Pareto optimality. We compute network-based travel times, i.e., weighted shortest paths, but build on previous work from Steinacker *et al.* (2022) by taking a demand-driven view on street network goodness by incorporating an origin-destination (OD) matrix. Let Ω be a set of OD pairs where the origin and destination are nodes in the graph, $\Omega = \{(u_1, v_1), (u_2, v_2), \dots\}$. Ω is derived from travel surveys, GPS trajectories, or bike sharing data. Let $t^c(e)$ and $t^b(e)$ be the travel times by car or bike along edge e, and let $P^c(u, v)$ and $P^b(u, v)$ be the edge set of a weighted shortest path from u to v, based on the edge weights t^c and t^b respectively. The Pareto frontier is computed to trade-off the goodness of the car network, $\mathcal{T}(G^c)$, with the goodness of the bike network $\mathcal{T}(G^b)$, defined as

$$\mathcal{T}(G^c) = \sum_{(u,v)\in\Omega} \Big(\sum_{e\in P^c(u,v)} t^c(e) \Big) \qquad \mathcal{T}(G^b) = \sum_{(u,v)\in\Omega} \Big(\sum_{e\in P^b(u,v)} t^b(e) \Big) \tag{2}$$

The edge-wise travel times $t_b(e)$ and $t_c(e)$ are set based on the lane's length $d(e)$ in km, its gradient $\delta(e)$ in % and its speed limit $\theta(e)$ in km/h, which are available from Open Street Map (OSM) data. The car travel times are simply set $t^c(e) = \frac{d(e)}{\theta(e)}$ since the gradient does not have a strong impact on the car speed in urban areas. Bike speed is estimated as 21.6km/h on flat lanes (excluding acceleration and breaking), with an increase of 0.86km/h per negative percent gradient (downhill acceleration) and an increase of 1.44km/h per positive percent gradient (uphill) (Parkin and Rotheram, 2010):

$$t^b(e) = \frac{d(e)}{v^b(e)} = \begin{cases} \frac{d(e)}{max\{1, \quad 21.6-1.44\cdot\delta(e)\}} & \text{if } \delta(e) > 0 \text{ (uphill)} \\ \frac{d(e)}{21.6-0.86\cdot\delta(e)} & \text{else (downhill or flat).} \end{cases} \tag{3}$$

However, a crucial component of cycling is safety, as a lack of dedicated infrastructure prevents people from cycling. This can be expressed in terms of a *perceived* bike travel time that amends the physical travel time with a psychological component. The perceived bike travel time is computed from the actual travel time, $t^b(e)$, by penalizing the discomfort of cycling on car lanes. The penalty is zero for edges equipped with dedicated bike lanes.

To set the penalty for roads without proper bike lanes, we follow the study by Meister *et al.* (2023), who recently studied route choices of cyclists in Zurich and found a "value of distance" of -0.66 for bike lanes and -0.36 for bike paths. Taking the average of both, we assume that the perceived distance is approximately halved on proper bike lanes, and we thus set perceived travel time for cycling on car lanes to $t^\beta(e) = 2 \cdot t^b(e)$.

3.2 A linear programming approach towards solving the BNAP

The goal of minimizing both bike and car travel times gives rise to an integer programming (IP) formulation with a multi-criteria objective function. To express travel times in an IP, we first revisit a flow formulation of the all-pairs shortest path problem:

$$\min \sum_{s,t \in V} \sum_{e \in E} f_{s,t,e} \cdot t(e) \tag{4}$$

$$s.t. \quad \forall v \in V, \quad \forall s,t \in V: \quad \sum_{e \in \delta^+(v)} f_{s,t,e} - \sum_{e \in \delta^-(v)} f_{s,t,e} = \begin{cases} -1 & \text{if } v = t \\ 1 & \text{if } v = s \\ 0 & \text{else,} \end{cases} \tag{5}$$

where $t \colon E \to \mathbb{R}_{\geq 0}$ encodes the travel times along the edges, and $f_{s,t,e}$ is the flow allocated on edge e for the path from s to t. Additionally $\delta^+(v)$ denotes the set of outgoing edges of node v and $\delta^-(v)$ its incoming edges. The flow constraints, together with the integer constraint $f_{s,t,e} \in \mathbb{Z}$, guarantee that there is a flow of value 1 between every (s,t)-pair, corresponding to a path. The objective computes the total travel time along each path, leading to the all-pairs shortest path in G.

We adapt the objective and constraints to model the BNAP. First, the undirected street graph G is converted into a directed graph G' by replacing every edge $e = (u,v) \in E$ by a pair of reciprocal directed edges, $\overrightarrow{e} = (u,v)$ and $\overleftarrow{e} = (v,u)$. The construction of G' allows to optimize the capacity for bike and car lanes in both directions. The edge properties of $\overrightarrow{e}$ and $\overleftarrow{e}$ are inherited from e; i.e., $d(\overrightarrow{e}) = d(\overleftarrow{e}) = d(e)$, $\theta(\overrightarrow{e}) = \theta(\overleftarrow{e}) = \theta(e)$, $\delta(\overrightarrow{e}) = -\delta(\overleftarrow{e})$ and $\lambda(\overrightarrow{e}) = \lambda(\overleftarrow{e}) = \lambda(e)$.

In the following, the adaptation of this formulation to the BNAP is explained in detail. An overview of all variables is given in Table 1 (Appendix A).

3.2.1 Introducing bike, car and shared flow

In contrast to the general shortest path formulation, we distinguish between the car flow f^c and the bike flow f^b along every edge, and frame the objective as a multi-criteria optimization problem. A weight γ can be set by the user to specify the desired importance of the car travel times relative to the bike travel times, resulting in the following preliminary objective function $\min \sum_{e \in E'} \sum_{(s,t) \in \Omega} f^b_{s,t,e} t^b(e) + \gamma \cdot f^c_{s,t,e} t^c(e)$ Varying γ yields Pareto-optimal solutions; i.e., scenarios where improving the bike network increases the car travel times and the other way around. Importantly, we compute the travel times over the s-t-pairs in an OD-matrix Ω, as motivated in subsection 3.1, instead of considering all node pairs. Additionally considering an OD-matrix reduces the runtime significantly. Specifically, with $n = |V|$ and $m = |E'|$, the all-pairs-shortest path formulation requires to optimize $n^2 m$ bike and car flow variables, since n^2 s-t-paths are considered, each with one variable per edge. This reduces to $|\Omega| m$ variables with the demand-driven approach. A drawback of this approach is that the resulting car network is not guaranteed to be strongly connected. To alleviate this problem, we add auxiliary pairs to Ω that form a chain of all vertices $((v_1, v_2), (v_2, v_3), \ldots, (v_{n-1}, v_n), (v_n, v_1))$.

Furthermore, the car flow constraints are set as in the all-pairs shortest-path formulation:

$$\forall v \in V, \quad \forall (s,t) \in \Omega: \quad \sum_{e \in \delta^+(v)} f^c_{s,t,e} - \sum_{e \in \delta^-(v)} f^c_{s,t,e} = \begin{cases} -1 & \text{if } v = t \\ 1 & \text{if } v = s \\ 0 & \text{else.} \end{cases} \tag{6}$$

However, this constraint is unsuitable for bike flow, since simultaneously guaranteeing bike and car flow along each path in Ω is oftentimes infeasible in a real-world street network. Thus, we introduce the concept of *shared flow*, denoted f^β, representing that cyclists can also use car lanes. Equation 7 expresses that a bike paths between all pairs of nodes s, t is only required with a combination of f^b (bike-on-bike-lane flow) and shared flow f^β (bike-on-car-lane flow).

$$\forall v \in V, \quad \forall (s,t) \in \Omega: \quad \sum_{e \in \delta^+(v)} (f^b_{s,t,e} + f^\beta_{s,t,e}) - \sum_{e \in \delta^-(v)} (f^b_{s,t,e} + f^\beta_{s,t,e}) = \begin{cases} -1 & \text{if } v = t \\ 1 & \text{if } v = s \\ 0 & \text{else.} \end{cases} \tag{7}$$

3.2.2 Constraining the space for bike and car lanes

Since the space on urban streets is limited, the goal of our approach is to decide which space to allocate to car and bike travel respectively. We model the space limitation with bike and car capacities, denoted λ^b_e and λ^c_e. The bike and car flow are constrained by the capacity: $\forall (s,t) \in \Omega, \forall e \in E': \quad f^c_{s,t,e} \leq \lambda^c_e \quad f^b_{s,t,e} \leq \lambda^b_e$

In turn, λ^b_e or λ^c_e are bounded by the total street capacity λ_e, which is given by the number of available lanes or the street width. Thus, the final constraint is in alignment with the

BNAP definition (Equation 1): $\forall e = (u, v) \in E' : \lambda^c_{(u,v)} + \lambda^c_{(v,u)} + 0.5(\lambda^b_{(u,v)} + \lambda^b_{(v,u)}) \leq \lambda(e)$. When λ_e is the number of lanes, it is useful to require bidirectional bike lanes ($\forall e \in E'$: $\lambda^b_{\overrightarrow{e}} = \lambda^b_{\overleftarrow{e}}$) to take up all available space. Notably, the shared flow f^β is not constrained. Instead, we penalize f^β, the undesired bike traffic on car lanes, in the objective function based on the higher perceived travel time $t^\beta(e)$. The objective function thus becomes

$$\min \sum_{(u,v)\in E'} \sum_{(s,t)\in\Omega} f^b_{s,t,e} t^b(e) + f^\beta_{s,t,e} t^\beta(e) + \gamma \cdot f^c_{s,t,e} t^c(e).$$

Solving the problem with the proposed constraints and objective yields the optimal bike and car capacities per street and direction, which can be interpreted as the number of lanes or the street width to allocate per transport mode.

3.3 Linear relaxation

The solution of the provided problem formulation can only be translated into bike and car lane-allocations if the problem is solved as an IP. However, the number of variables to optimize remains large even with the proposed relaxations, rendering IP computationally prohibitive for real-world instances. Therefore, we solve the problem as a linear program, resulting in fractional flow values for the bike and car capacities (λ^b_e, λ^c_e). We propose an iterative process of rounding the capacities and re-computing the optimal solution. For an algorithmic description of the post-processing, see also algorithm 1 (Appendix B). Let Λ be a set of the indices and values of all *fixed* capacities, $\Lambda = \{(e, i, \lambda^i_e) \mid i \in \{b, c\},\ e \in E\}$. Initially, Λ is empty ($\Lambda = \varnothing$), or Λ corresponds to a set of lanes that are fixed due to real-world constraints such as compulsory car lanes that are used by bus services. In each iteration, the LP is solved subject to the fixed bike capacities Λ, yielding the optimal capacities λ^*, where $\forall (i, e, \lambda^i_e) \in \Lambda : \quad \lambda^i_e = \lambda^{*,i}_e$ (i.e., the fixed capacities remain unchanged). The algorithm then rounds up the k largest *bike* capacity values and fixes them. Afterwards the solution is recomputed to optimize the remaining capacities. Before fixing a bike lane, it is ensured that the remaining car network remains strongly connected. Each iteration results in a feasible graph division into car and bike network, assuming all lanes aside from the fixed bike lanes are car lanes.

4 Experiments

We test the presented evaluation and optimization framework on real data from Zurich (Switzerland), Cambridge MA(US) and Chicago (US). From each city, two districts were selected, and their street network was extracted from Open Street Map (OSM) and pre-processed with the SNMan Python library[1] by Ballo and Axhausen (2024) (see Appendix C). An overview of the six instances is shown in Figure 1 (for numerical

[1]https://github.com/lukasballo/snman

properties, see Table 2 in Appendix C). The OD-matrices Ω are derived from public bike sharing data (Chicago & Cambridge) or census data (Zurich); see Appendix C for details. In the following, we report the travel times over OD-paths in Ω if not denoted otherwise. All tests were executed on a standard machine with 16 GB RAM, using a Gurobi solver. The source code is available at `https://github.com/mie-lab/bike_lane_optimization`.

Figure 1: Network layout of the real instances used to test our algorithm

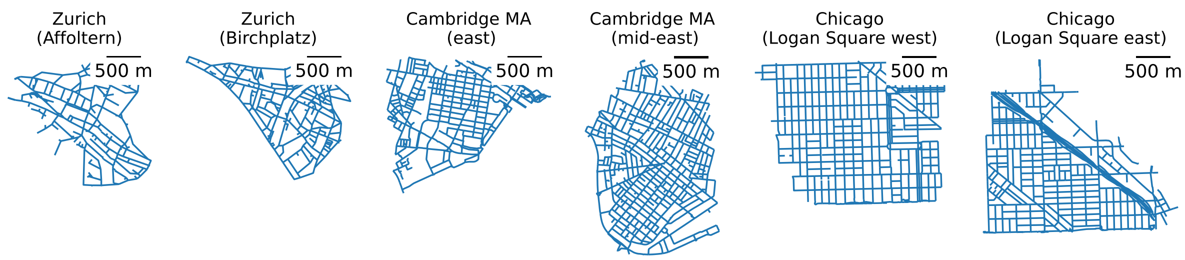

To benchmark the novel algorithm with respect to prior work, we compare to three heuristic methods based on the betweenness centrality. Since Steinacker *et al.* (2022) do not consider the impact of bike lanes on the car network, we adapt their algorithm in three ways. First, we assume a full network of bike *priority* lanes, where cars are restricted to 10km/h, and iteratively assign edges to the car network (*betweenness-top-down*). Secondly, starting from a car network without any bike lanes, edges are iteratively allocated for cycling, starting with the lanes with lowest betweenness centrality in the *car* network (*betweenness-bottom-up (car)*). Third, the same process is repeated but starting the reallocation-process with lanes that have high betweenness centrality with respect to *bike* travel (*betweenness-bottom-up (bike)*). For implementation details, see Appendix D.

Figure 2 presents the Pareto frontier for each method and instance, visualizing the achieved trade-off between bike and car travel time. Each point on the Pareto frontier is one bike network. For example, for Zurich-Affoltern our algorithm yields a network where the perceived bike travel time is decreased by 40% while increasing the car travel time only by 17%. The bike travel time can be reduced by at most 50% due to the setting of $t^{\beta}(e) = t^{b}(e)$. In five instances, the Pareto frontier of our optimization approach dominates over the heuristic solutions. The bike-focused bottom-up method usually yields better networks when many bike lanes are allocated, whereas the car-focused bottom-up method yields better solutions when few bike lanes are allocated (see intersections of green and yellow lines for Cambridge MAand Chicago). The framework thus also provides recommendations about the use cases for the respective methods.

Figure 2: Pareto optimality of bike networks. Top: Algorithms are compared by their Pareto frontier. In five out of six instances, our linear programming approach outperforms methods based on the betweenness centrality. Bottom: Each point on the Pareto frontiers (top) corresponds to one plausible street network. Three examples in Cambridge MAare shown, where the bike networks differ dependent on the planning method and the number of allocated bike lanes. This is also reflected in the distance of shortest paths, where the existence of dedicated cycling infrastructure is rewarded in the perceived bike travel time.

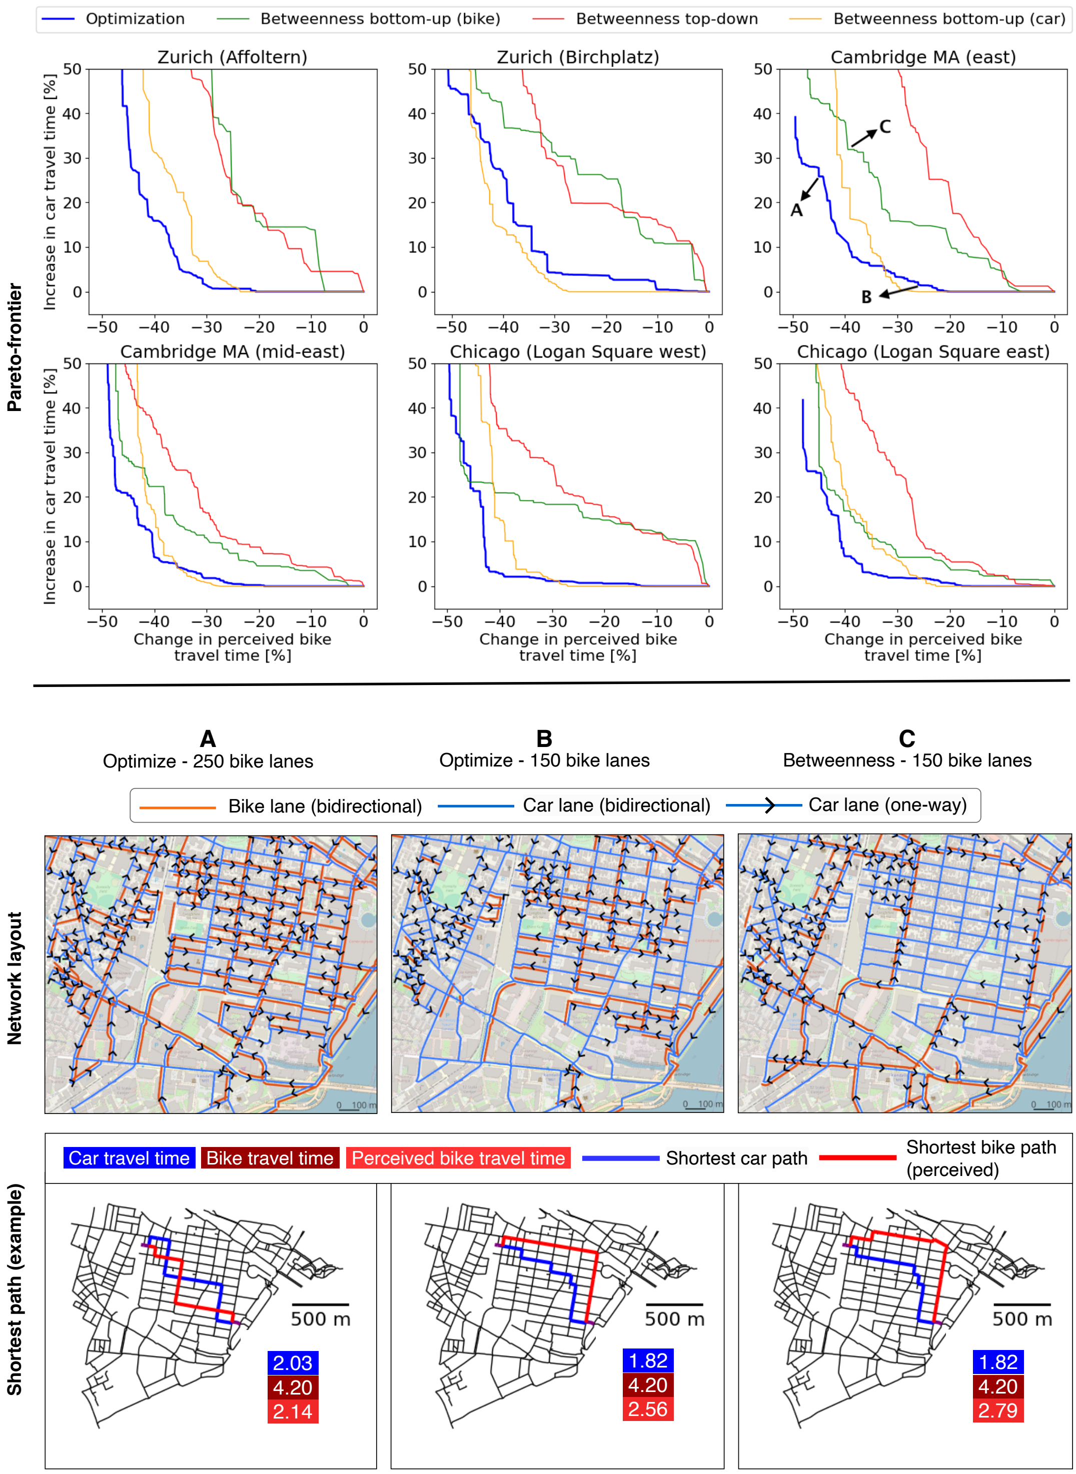

Furthermore, Figure 2 (bottom) illustrates three resulting bike networks for Cambridge MA(east), varying due to the planning method and the number of allocated bike lanes. The heuristic method (C) places a greater number of bike lanes on main roads compared to our optimization approach (B). When allocating more bike lanes (250 instead of 150) with the optimization approach, the car network transitions into a complex one-way system. This change leads to increased car travel times, as depicted in Figure 2 (bottom), where travel times increase from 1.82 to 2.03 for a specific route. Although the overall bike travel time remains consistent across all algorithms - since cyclists have access to all roads - the perceived bike travel time varies significantly, and it is considerably longer when using the betweenness algorithm compared to the optimization method, highlighting the impact of planning approaches on travel efficiency.

5 Discussion

Our study introduces a novel perspective on bike network planning by prioritizing the trade-off between car and bike travel times through Pareto optimality. The aim of minimizing the impact of bike lanes on other modes is formalized in the bike lane allocation problem (BNAP) and addressed with our IP formulation that yields Pareto-optimal solutions. To enable the application to real-world scenarios, we developed several relaxations and post-processing schemes, while also integrating key innovations such as demand-driven aspects and the assessment of perceived bike travel times.

Addressing actual traffic flow remains a significant hurdle. Accurately modeling traffic flow typically requires sophisticated simulators, making it challenging to employ optimization algorithms without resorting to a bi-level formulation. However, ignoring traffic flow has certain effects on the optimal solution. For example, our algorithm is predisposed to allocate one lane of any double-lane road to bicycles, not accounting for the impact on car travel times due to reduced road width.

Future work should thus focus on accurately modeling traffic flow and on adapting our algorithm to account for factors critical to real-world traffic dynamics, such as intersection layouts. Additionally, there are straightforward extensions, such as incorporating parking space allocation constraints or incentivizing bike lanes near green spaces, that could further refine the model. To make our algorithm more accessible, developing a user interface that allows parameter adjustments through sliders and visualizes the resultant street networks is crucial. Such a tool would empower urban planners with limited technical expertise to leverage our algorithm for informed decision-making.

6 References

Ballo, L. and K. W. Axhausen (2024) Modeling sustainable mobility futures using an automated process of road space reallocation in urban street networks: A case study in Zurich, *103rd Annual Meeting of the Transportation Research Board, Washington DC*.

Ballo, L., L. M. de Freitas, A. Meister and K. W. Axhausen (2023) The E-Bike City as a radical shift toward zero-emission transport: Sustainable? Equitable? Desirable?, *Journal of Transport Geography*, **111**.

Bao, J., T. He, S. Ruan, Y. Li and Y. Zheng (2017) Planning bike lanes based on sharing-bikes' trajectories, paper presented at the *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1377–1386.

Boeing, G. (2017) OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks, *Computers, Environment and Urban Systems*, **65**, 126–139.

Brandes, U. (2008) On variants of shortest-path betweenness centrality and their generic computation, *Social networks*, **30** (2) 136–145.

Buehler, R. and J. Dill (2016) Bikeway networks: A review of effects on cycling, *Transport reviews*, **36** (1) 9–27.

Burke, C. M. and D. M. Scott (2016) The space race: A framework to evaluate the potential travel-time impacts of reallocating road space to bicycle facilities, *Journal of Transport Geography*, **56**, 110–119.

Duthie, J. and A. Unnikrishnan (2014) Optimization framework for bicycle network design, *Journal of Transportation engineering*, **140** (7) 04014028.

EDGAR/JRC (2022) Distribution of carbon dioxide emissions worldwide in 2022, by sector, *Statista*.

Fraser, S. D. and K. Lock (2011) Cycling for transport and public health: A systematic review of the effect of the environment on cycling, *The European Journal of Public Health*, **21** (6) 738–743.

Gerike, R., S. Weikl, C. Koszowski and K. Bogenberger (2022) Network level design for cycling, in *Advances in Transport Policy and Planning*, vol. 10, 77–109, Elsevier.

Grisiute, A., N. Wiedemann, H. Martin and M. Raubal (2023) Ontology-based approach for harmonizing metrics in bike network evaluations, paper presented at the *Cycling Research Board Annual Meeting 2023. Book of Abstract*, 113–116.

Lin, J.-J. and C.-J. Yu (2013) A bikeway network design model for urban areas, *Transportation*, **40**, 45–68.

Liu, S., Z.-J. M. Shen and X. Ji (2022) Urban bike lane planning with bike trajectories: Models, algorithms, and a real-world case study, *Manufacturing & Service Operations Management*, **24** (5) 2500–2515.

Mauttone, A., G. Mercadante, M. Rabaza and F. Toledo (2017) Bicycle network design: model and solution algorithm, *Transportation Research Procedia*, **27**, 969–976.

Meister, A., M. Felder, B. Schmid and K. W. Axhausen (2023) Route choice modeling for cyclists on urban networks, *Transportation Research Part A: Policy and Practice*, **173**, 103723.

Mueller, N., D. Rojas-Rueda, M. Salmon, D. Martinez, A. Ambros, C. Brand, A. De Nazelle, E. Dons, M. Gaupp-Berghausen, R. Gerike *et al.* (2018) Health impact assessment of cycling network expansions in european cities, *Preventive medicine*, **109**, 62–70.

Natera Orozco, L. G., F. Battiston, G. Iñiguez and M. Szell (2020) Data-driven strategies for optimal bicycle network growth, *Royal society open science*, **7** (12) 201130.

Oja, P., S. Titze, A. Bauman, B. De Geus, P. Krenn, B. Reger-Nash and T. Kohlberger (2011) Health benefits of cycling: A systematic review, *Scandinavian journal of medicine & science in sports*, **21** (4) 496–509.

Olmos, L. E., M. S. Tadeo, D. Vlachogiannis, F. Alhasoun, X. E. Alegre, C. Ochoa, F. Targa and M. C. González (2020) A data science framework for planning the growth of bicycle infrastructures, *Transportation research part C: emerging technologies*, **115**, 102640.

Parkin, J. (2018) *Designing for cycle traffic: International principles and practice*, ICE Institution of Civil Engineers Publishing.

Parkin, J. and J. Rotheram (2010) Design speeds and acceleration characteristics of bicycle traffic for use in planning, design and appraisal, *Transport Policy*, **17** (5) 335–341.

Paulsen, M. and J. Rich (2023) Societally optimal expansion of bicycle networks, *Transportation Research Part B: Methodological*, **174**, 102778.

Schoner, J. E. and D. M. Levinson (2014) The missing link: Bicycle infrastructure networks and ridership in 74 US cities, *Transportation*, **41**, 1187–1204.

Steinacker, C., D.-M. Storch, M. Timme and M. Schröder (2022) Demand-driven design of bicycle infrastructure networks for improved urban bikeability, *Nature Computational Science*, 1–10.

Szell, M., S. Mimar, T. Perlman, G. Ghoshal and R. Sinatra (2022) Growing urban bicycle networks, *Scientific Reports*, **12** (1) 6765.

Ta, N., H. Li, Y. Chai and J. Wu (2021) The impact of green space exposure on satisfaction with active travel trips, *Transportation Research Part D: Transport and Environment*, **99**, 103022.

Weikl, S. and P. Mayer (2023) Data-driven quality assessment of cycling networks, *Frontiers in Future Transportation*, **4**, 1127742.

Yan, X., T. Wang, X. Ye, J. Chen, Z. Yang and H. Bai (2018) Recommended widths for separated bicycle lanes considering abreast riding and overtaking, *Sustainability*, **10** (9) 3127.

A Overview of all variables

Table 1: List of variables and explanation

Variable	Description
Optimized variables	
$f^c_{s,t,e}$	Car flow on car lane along edge e for the path from s to t
$f^b_{s,t,e}$	Bike flow on bike lane along edge e for the path from s to t
$f^\beta_{s,t,e}$	Bike flow on car lane (shared lane) along edge e for the path from s to t
λ^c_e	Capacity allocated for cars on edge e
λ^b_e	Capacity allocated for bikes on edge e
Inputs	
$G = (V, E)$	Street graph with one edge per street
$G' = (V, E')$	Auxiliary street graph with a pair of reciprocal directed edges per street
n	Number of nodes
m	Number of edges
Ω	Set of considered origin-destination pairs, given as pairs of node $(u, v) \in V^2$
$t^b(e)$	Travel time by bike along edge e
$t^c(e)$	Travel time by car along edge e
$t^\beta(e)$	Perceived travel time for cycling on a car lane along edge e
γ	Weighting (desired importance) of the car travel time
λ_e	Overall given capacity of edge e

B Post-processing algorithm

Algorithm 1: Post-processing scheme to round the LP solution

```
Input: Street network G with edge attributes d(e), δ(e), θ(e)
Input: Set of fixed capacities Λ
Input: Ω, ω, γ, k
i = 0
repeat
   λ* = LP(G, Λ, Ω, ω, γ)
   Sort λ*_b ;                          // Round edges with largest bike capacity
   for k iterations do
      ê = edge with largest λ*_b that is not in Λ yet
      if G remains strongly connected then
         Λ = Λ ∪ {(ê, b, 1)} ;                           // Fix as bike lane
      else
         Λ = Λ ∪ {(ê, c, 1)} ;                            // Fix as car lane
      end if
   end for
   i = i + 1
until |Λ| = 2m ;          // Until all edges are fixed as bike or car lanes
```

C Preprocessing of real network data

The SNMan library builds up on the `networkx` and `OSMnx` (Boeing, 2017) packages and first constructs a *street* graph with one node per intersection and one edge per street. It further provides functionality to convert the street graph into the *lane graph*, a directed multigraph. Each node in the lane graph is defined by geographic coordinates and elevation, and the edges are enriched with attributes for their distance, speed limit and the type of lane based on available OSM data, allowing to derive $\delta(e)$, $t^c(e)$ and $t^b(e)$.

An origin-destination Ω expressing real-world travel demand is derived from public bike sharing or census data. For districts in Zurich, we take all trips in the Mobility Microcensus from Switzerland that intersect with the district region. The Mobility and Transport Microcensus is a statistical survey on travel behaviour that is published by the Federal Office for Spatial Development. Overall, it contains travel survey data for more then 57k participants. After intersecting the origin-destination-lines with the city district, 1061 and 508 trips remain for Birchplatz and Affoltern respectively. The trip origin and destination are matched to nodes in the graph by selecting the node closest to their geographic coordinates, yielding 498 and 267 unique OD-pairs respectively (see Table 2).

On the other hand, to the best of our knowledge there is no public travel data for Chicago or Cambridge. Instead, we utilize public data from their respective local bike sharing services. It is worth noting that this biases the OD-data to cycling movement; however, it can be assumed that the data still reflect typical mobility behaviour within the district, irrespective of the transport mode. For Chicago, data is available from Divvy bike sharing[2], whereas in Cambridge, the operating service is Bluebikes[3]. In both cases, we download and merge all trip data from the whole year of 2023 and only filter out stations that are marked as charging or maintenance stations. The origin-destination data is intersected with the respective district based on their beeline connection, and the resulting subset of OD-pairs is mapped to graph nodes by their geographic locations, yielding a set of node-based OD-pairs. Due to the large number of resulting pairs, we select only the OD-pairs that collectively account for 75% of the trips from 2023, resulting in the counts listed in Table 2.

D Heuristic baselines

To demonstrate our evaluation framework based on Pareto optimality, we compare our LP algorithm to three heuristic methods inspired by previous work. Steinacker *et al.*

[2] https://divvybikes.com/system-data
[3] https://bluebikes.com/system-data

Table 2: Overview of real-world instances. The runtime for solving the LP one time is given as "Runtime optimization", in contrast to the runtime for computing the whole Pareto frontier ("Runtime Pareto").

	Zurich		Cambridge		Chicago	
	Affoltern	Birchplatz	east	mid-east	L.S. west	L.S. east
Nodes	213	301	283	504	345	371
Streets	290	431	506	775	601	603
Lanes	535	799	883	1253	934	1126
OD-paths	219	383	456	677	175	281
OD-paths extended	430	684	738	1180	520	648
Runtime optim. [min]	4	36	20	146	26	38
Runtime Pareto [h]	0.28	4.3	2.74	7.91	5.35	7.07

(2022) proposed to generate a sequence of bike networks, starting from a network with bike lanes at every edge, and iteratively removing bike lanes based on their betweenness centrality, an index measuring how frequently an edge is part of a shortest path (Brandes, 2008). In their work, the betweenness centrality is computed with respect to the shortest paths of an OD-matrix derived from the pickups and drop-offs in a bike sharing system. Our work builds up on this demand-driven approach; however, Steinacker *et al.* (2022) ignore the effect of new bike lanes on other traffic, which impedes the comparability of the networks resulting from their and our algorithm. Nevertheless, we implemented three baseline approaches that are based on their approach and that utilize the betweenness centrality measure. For the first one, we follow (Steinacker *et al.*, 2022) closely and start from a network where all streets are bike lanes. To express the negative impact of proper bike lanes on the car network, we exploit the concept of *bike priority lanes*, where cars are required to give priority to bikes and are slowed down accordingly. Specifically, it is assumed that cars can drive 10 km/h on bike lanes. The initial network is thus a full bike lane network with a car speed of 10km/h throughout the city. As in (Steinacker *et al.*, 2022), edges are iteratively removed from the network (and designated as proper car lanes), starting from the edge with lowest betweenness centrality with respect to the shortest bike travel times. The same OD-matrix and (perceived) travel times as for our optimization approach are used to ensure comparability. Since this method starts from a full bike network, we call this approach *betweenness-top-down*.

In contrast, for the second and third baseline, we start from a network with car lanes only and add bike lanes iteratively, termed *betweenness-bottom-up* in the following. There are two ways to re-assign lanes to cycling: 1) a car-prioritizing approach, where the first edges to be converted to bike lanes are the ones with *lowest* betweenness centrality, computed with respect to the *car* travel time (*betweenness-bottom-up (car)*), and 2) a bike-prioritizing approach, where the edges with *highest* betweenness centrality with respect to the *bike* travel time are converted first (*betweenness-bottom-up (bike)*). In both cases, we iteratively select the edge with the lowest/highest betweenness centrality computed on the OD-matrix, and convert this edge into a bidirectional bike lane, setting the car travel time along this edge to ∞ as in our optimization approach. In summary, we compare our approach to three strong heuristic methods based on the betweenness centrality, where the first one (top-down) is designed to be as similar as possible to (Steinacker *et al.*, 2022) while the second and third one (bottom-up) are constructed for better comparability to our method. The method *betweenness-bottom-up (car)* resembles the approach taken in Ballo and Axhausen (2024).