Eindhoven University of Technology

MASTER

Demand-Informed Optimization of Cycle Service Point Locations in Urban Areas

Hiemink, M.A.

Award date:
2026

Link to publication

Disclaimer
This document contains a student thesis (bachelor's or master's), as authored by a student at Eindhoven University of Technology. Student theses are made available in the TU/e repository upon obtaining the required degree. The grade received is not published on the document as presented in the repository. The required complexity or quality of research of student theses may vary by program, and the required minimum study period may vary in duration.

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 25. Nov. 2026

EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of the Built Environment
Urban Planning and Transportation Research Group

Demand-Informed Optimization of Cycle Service Point Locations in Urban Areas

Master's Thesis

M.A. Hiemink
Student number: 1453297

Master's programme Architecture, Building and Planning
Track: Urban Systems and Real Estate
Graduation Project (7Z45M0)
45 ECTS

Supervisors:
Ir. D.T.L. Andreoli
Dr. F. Liao
Ir. S.J. Buningh

September 25, 2026

This graduation thesis is publicly available.

This thesis has been carried out in accordance with the rules of the TU/e Code of Scientific Integrity.

Preface

This thesis marks the completion of my Master of Science in Architecture, Building and Planning at Eindhoven University of Technology, within the Urban Systems and Real Estate track. It brings together the spatial, behavioral and computational strands of the program in a single question, where cycling commuters can best be supported, and it has been the most demanding and the most rewarding work of my studies.

The subject of this thesis originates with BAM Infra, whose Cycle Service Point concept motivated the research question. I am grateful for the opportunity to work on a problem with such a direct practical grounding.

I would like to thank the members of my graduation committee for their guidance throughout this project. I am grateful to Ir. D.T.L. Andreoli, my day-to-day supervisor, for his availability and assistance in making difficult decisions. I thank Dr. F. Liao, my second supervisor, for assisting me with the more technical aspects of the models I employed. I am grateful to Ir. S.J. Buningh, my external supervisor whose business experience provided practical guidance.

Finally, I thank my family for their never-ending support, my friends for the necessary distractions, and my fellow students for the discussions and companionship that made these years what they were. This work would not have been possible without them.

Mats Hiemink
Sunday 20th September, 2026

Abstract

Cycling is central to sustainable urban mobility, yet the facilities that support commuter cycling are being planned without a systematic, demand-informed method for deciding where they should be placed. The Cycle Service Point, a modular node offering secure parking, minor repair and rest, is one such facility. This thesis develops a generalizable framework that combines an agent-based cycling-demand model with a flow-capturing siting model, and applies it to the metropolitan region of Eindhoven. What sets the framework apart is the demand it sites against. Instead of placing facilities near where people live, it places them on the street-by-street routes that simulated cyclists actually ride, so that each candidate location is judged by how many passing cyclists it can reach. Phase 1 builds that demand in four steps. A regression estimates how many cycling trips each zone produces, the 2022 Dutch national travel survey is expanded into a synthetic population in which every simulated cyclist carries a full-day travel diary, a gravity model gives each trip a destination, and the agent-based transport simulator MATSim routes the trips over the cycle network. The result is an hourly, street-level picture of 644,904 weekday bicycle trips made by 250,559 simulated cyclists. It reproduces bicycle counting-station volumes with a Nash–Sutcliffe efficiency of 0.518 and a mean scalable quality value of 0.905, and 78.1% of hourly station comparisons fall within the tolerance normally applied to traffic counts. This falls short of the strictest guideline for motor-traffic models but is an acceptable fit for a strategic, region-scale cycling model: the model reproduces where and how much cycling occurs across the region, although less precisely at individual stations and hours. Phase 2 scores candidate locations on eight open-data criteria and selects the set that intercepts the most cycling flow, using a greedy algorithm with a guaranteed bound on its distance from the best possible solution. Five service points capture 8.5% of peak commuter flow, concentrated in Eindhoven and Helmond. Because demand is resolved at the level of street links, the defensible output is a ranked set of corridor stretches rather than exact coordinates. Every input is national open data, so the framework can be used for any other Dutch urban region by substitution.

Contents

Contents iv

List of Figures vi

List of Tables viii

List of Formulas ix

1 Introduction 1
- 1.1 Background . . . 1
- 1.2 Problem Statement . . . 2
- 1.3 Research Objectives . . . 3
- 1.4 Scope . . . 4
- 1.5 Thesis Structure . . . 4

2 Literature Study 6
- 2.1 Cycling Behavior and Mode Choice . . . 6
- 2.2 Cycling Infrastructure and Commuter Support Facilities . . . 9
- 2.3 Mobility Hubs and Multimodal Transport Integration . . . 9
- 2.4 Trip Generation Regression for Cycling Demand Scaling . . . 12
- 2.5 Agent-Based Modeling and MATSim for Transport Simulation . . . 13
- 2.6 Optimized Location Search for Cycling Infrastructure . . . 15
- 2.7 Summary and Identified Research Gap . . . 17

3 Methodology 19
- 3.1 Introduction . . . 19
- 3.2 Framework Overview . . . 20
- 3.3 Data Sources . . . 23
- 3.4 Phase 1: Agent-Based Cycling Demand Model . . . 23
- 3.5 Phase 2: Demand-Informed Siting of Cycle Service Points . . . 38
- 3.6 Ethical Considerations . . . 49
- 3.7 Conclusion . . . 49

4 Phase 1 Results: A Validated Cycling-Demand Surface for the MRE 51
- 4.1 Trip Generation . . . 51
- 4.2 Synthetic Population and Demand Calibration . . . 58
- 4.3 Trip Distribution . . . 60
- 4.4 Network Assignment and the Demand Surface . . . 63
- 4.5 Model Validation . . . 68
- 4.6 Conclusion . . . 70

5 Phase 2 Results: Siting of Cycle Service Points **72**
5.1 The Demand Object and the Two Models . . . 72
5.2 Candidates and the Capturable-Flow Surface . . . 74
5.3 The Eight Factors and Their Overlap . . . 77
5.4 Coverage . . . 83
5.5 Site Structure and Cross-Model Agreement . . . 83
5.6 Robustness . . . 85
5.7 Conclusion . . . 88

6 Conclusions and Discussion **89**
6.1 Conclusions . . . 89
6.2 Practical and Policy Relevance . . . 90
6.3 Discussion of the Model . . . 91
6.4 Limitations and Future Work . . . 94

Bibliography **96**

A Supporting Tables **101**
A.1 Observed Cycling Speeds by Infrastructure Type . . . 101
A.2 Trip Generation Model Diagnostics . . . 102
A.3 Population Synthesis Configuration . . . 103
A.4 Facility Activity-Type Mapping . . . 104
A.5 Distance-Decay Calibration . . . 105
A.6 MATSim Simulation Control Parameters . . . 106
A.7 Additional Validation Metrics . . . 107
A.8 Per-Factor Suitability Surfaces . . . 108

List of Figures

3.1 The Metropolitan Region Eindhoven (MRE) (study area). 20
3.2 Framework of the thesis 22
3.3 Filtering flow of the 2022 ODiN sample 26

4.1 Predictor correlation 52
4.2 ODiN survey coverage 53
4.3 Observed vs predicted zonal cycling 55
4.4 Person expansion factors by zone 57
4.5 Person expansion-factor distribution 57
4.6 Cycling trips per inhabitant 58
4.7 Donor diversity per zone 59
4.8 Own-zone donor share 59
4.9 Synthetic agents by zone 60
4.10 Assigned distances by purpose 61
4.11 Assigned vs survey distances 62
4.12 MATSim score convergence 63
4.13 Concentration of demand over links 64
4.14 Cycling demand field 65
4.15 Busiest corridors 66
4.16 Hourly cycling profile 67
4.17 Executed trip lengths 68
4.18 Simulated vs observed station-hour volumes 69
4.19 GEH distribution 70
4.20 Observed vs simulated by hour 70

5.1 Peak cycling demand field 73
5.2 All-day cycling demand field 74
5.3 Capturable-flow surface 75
5.4 Surface resolution 76
5.5 Factor-overlap matrix 78
5.6 Per-factor surfaces 79
5.7 Betweenness siting 80
5.8 Accessibility siting 81
5.9 Route-junction siting 81
5.10 Transit-hub siting 81
5.11 Business-density siting 82
5.12 Urban-core siting 82
5.13 Land-use-mix siting 82
5.14 Cycle-network siting 83
5.15 Selected sites by model 84
5.16 Demand by municipality 85
5.17 Per-site close-ups 85

5.18 Coverage and shared sites under future demand 86

A.1 Route-junction surface 109
A.2 Betweenness surface 109
A.3 Land-use-mix surface 110
A.4 Urban-core surface 110
A.5 Business-density surface 111
A.6 Accessibility surface 111
A.7 Transit-hub surface 112
A.8 Cycle-network surface 112

List of Tables

3.1 Largest link exclusion classes in the bicycle-access filter . . . 28
3.2 Cycling attractiveness scores over the network . . . 28
3.3 Phase 2 input layers . . . 40
3.4 The five factor forms . . . 41
3.5 Factor forms, scales and value functions . . . 42
3.6 The two weight profiles . . . 45
3.7 Current scenario weights . . . 48

4.1 Candidate trip-generation predictors . . . 52
4.2 NB2 trip-generation model . . . 53
4.3 Trip-generation fit statistics . . . 56
4.4 Synthetic population against the ODiN weights . . . 59
4.5 Calibrated gravity distance-decay parameters . . . 61
4.6 Assigned demand by purpose . . . 62
4.7 Link-level cycling demand . . . 64
4.8 Validation against NDW counts . . . 69

5.1 Phase-2 demand and candidate scale . . . 73
5.2 Demand and candidate funnel . . . 75
5.3 The eight quality factors . . . 77
5.4 Factor value distributions . . . 77
5.5 Coverage from single-factor siting . . . 79
5.6 Coverage by model . . . 83
5.7 Selected sites and their rivals . . . 84
5.8 Coverage under current and future demand scenarios . . . 86
5.9 Weight sensitivity of coverage . . . 87
5.10 Structural sensitivity . . . 87

A.1 Observed cycling speeds by infrastructure type . . . 101
A.2 Candidate predictor multicollinearity . . . 102
A.3 Cross-validation of the trip generation model . . . 102
A.4 Population synthesis configuration . . . 103
A.5 Facility activity-type mapping . . . 104
A.6 MATSim simulation control parameters . . . 106
A.7 Per-station validation . . . 108

List of Formulas

3.1 Weekday population estimate from the ODiN person weight 26
3.2 Average-weekday trip control total from the ODiN trip weight 27
3.3 Attractiveness-based additional link disutility 28
3.4 NB2 conditional mean with a population exposure offset 30
3.5 NB2 conditional variance as a quadratic function of the mean 30
3.6 Donor selection weight for population synthesis 31
3.7 Own-zone advantage factor in donor selection 31
3.8 Gaussian similarity between the target zone and a donor's home zone 31
3.9 Attribute typicality of a candidate donor for the target zone 32
3.10 Number of synthetic agents to draw per zone under trip anchoring 32
3.11 Gravity attractiveness of a candidate destination facility 34
3.12 Destination choice probability over the candidate facility set 34
3.13 Initial distance-decay parameter from the observed mean trip distance 34
3.14 Conversion of observed network distance to a straight-line target 34
3.15 Candidate search radius capturing 99% of the decay mass 34
3.16 Charypar–Nagel plan utility with a bicycle infrastructure term 35
3.17 Relative variation of the mean executed score (convergence criterion) 36
3.18 GEH statistic for comparing simulated and observed volumes 37
3.19 Nash–Sutcliffe efficiency of simulated against observed volumes 37
3.20 Quality-weighted flow capture of a set of CSP positions 39
3.21 Linear distance-decay capture kernel 39
3.22 Multi-criteria quality of a candidate position 44
3.23 The CSP selection problem under a fixed budget 45
3.24 Online submodular upper bound on the optimum 46
3.25 CRITIC objective weights 47
3.26 Band-targeted future-demand multiplier 48
A.1 Scalable Quality Value for volume comparison 107
A.2 Kling–Gupta efficiency combining correlation, variance and bias 107

List of Abbreviations

ABM Agent-Based Modeling.
AIC Akaike Information Criterion.
BAG Basisregistratie Adressen en Gebouwen.
BBG Bestand Bodemgebruik.
CBS Statistics Netherlands.
COROP Coördinatie Commissie Regionaal OnderzoeksProgramma.
CSP Cycle Service Point.
DMRB Design Manual for Roads and Bridges.
FCLM Flow-Capturing Location Model.
GA Genetic Algorithm.
GEH Geoffrey E. Havers.
GIS Geographic Information System.
GPS Global Positioning System.
KGE Kling–Gupta efficiency.
LSCP Location Set Covering Problem.
MAE Mean Absolute Error.
MATSim Multi-Agent Transport Simulation.
MCA Multi-Criteria Analysis.
MCLP Maximal Covering Location Problem.
MRE Metropolitan Region Eindhoven.
NB2 Type-2 Negative Binomial.
NDW Nationaal Dataportaal Wegverkeer.
NSE Nash–Sutcliffe efficiency.
ODiN Onderweg in Nederland.
OLS Ordinary Least Squares.
OSM OpenStreetMap.
PC4 4-digit postcode zone.
PDOK Publieke Dienstverlening Op de Kaart.
QSim queue-based mobility simulation.
RMSE Root Mean Squared Error.
SQV Scalable Quality Value.
VIF Variance Inflation Factor.

Chapter 1

Introduction

The United Nations (2018) projects that cities will account for approximately two-thirds of global population growth by 2050, adding an estimated 2.5 billion urban residents within three decades. This demographic concentration will place unprecedented pressure on urban transport systems already operating at or near capacity, intensifying demand for movement within increasingly dense spatial configurations (Chang et al., 2017).

The dominant response to growing urban transport demand has historically been the expansion of car-centric infrastructure. This strategy increases vehicular throughput and has generated well-documented externalities. As Sheller and Urry (2000) argue, car-dependent urban design structurally marginalizes non-motorized users, producing spatial inequalities in which access to mobility itself becomes a marker of social inclusion. Concurrently, even in infrastructure environments designed to favor the automobile, congestion persists as one of the most consequential and politically salient challenges confronting urban planners (Barthelemy, 2016). These conditions collectively motivate a structural reorientation toward sustainable urban mobility alternatives.

Active mobility, such as cycling and walking, offers a sustainable and equitable alternative to car-dependent transport by reducing congestion, improving public health, and enhancing urban connectivity. By integrating well-designed infrastructure with supportive policies, cities can foster a shift toward healthier, more inclusive, and environmentally friendly transport behaviors (Hull and O'Holleran, 2014).

1.1 Background

Cycling is central to that reorientation. Research on active mobility and its influence on mode choice is growing, and dedicated cycling infrastructure is being built to accommodate cyclists and to offer an alternative to the car. Whether people choose to cycle, however, depends on how cycling compares with the alternatives for the trips they actually make.

For commuting in particular, the choice of mode is shaped by perceived convenience, cost, travel time and reliability. Cycling is often preferred when it offers shorter travel times, lower costs and greater flexibility than the car or public transport, while infrastructure quality, safety concerns and weather can deter it (Heinen et al., 2010). Targeted incentives can shift this balance. de Kruijf et al. (2018) shows that an incentive program moved car commuters to the bicycle, with an effect that persisted for commutes longer than 20 km, especially among e-bike users.

Cycling to work is favored by short distances, safe infrastructure and workplace facilities such as showers and secure parking, while weather, car ownership and cultural norms shape how feasible and appealing it is (Heinen et al., 2010). Policies that aim to increase cycling to work therefore have to address practical as well as perceptual barriers, so that cycling becomes a reliable option that commuters can build into their daily routine.

Several of these barriers concern the journey itself rather than the network it runs on. Inadequate provision for secure parking and minor maintenance, limited access to rest and refreshment along commuting corridors, and the perceived risk of bicycle theft are identified as deterrents, particularly for car-dependent commuters (Heinen et al., 2010; Fishman et al., 2012). Research on cycling infrastructure has long concentrated on lane design, separation from motor traffic and intersection safety, but support facilities receive growing attention. Secure storage, changing rooms and showers at the workplace raise the likelihood of cycling to work (Heinen et al., 2013), although the evidence that such facilities also increase how often people cycle is mixed (Heinen et al., 2010). These barriers are therefore open to targeted intervention, but the size of the effect depends on the setting.

The Cycle Service Point (CSP), a modular facility concept, responds to this barrier profile by combining secure parking, minor repair services and commuter amenities at nodes within the cycling network (Loendersloot, 2022). The concept is new. The first facility, the Fietsservicepunt Schiphol Noord in the Amsterdam region, went into construction at the end of 2025 (BAM Infra Nederland, 2025), and the concept itself has received little academic attention. Its expected benefit therefore rests on the wider evidence on support facilities rather than on evaluations of CSPs themselves, and that evidence points to location as a deciding factor: a facility can only lower barriers for the cyclists who pass it, and the mobility-hub literature names public indifference, hubs that go largely unused, as one of the main risks of implementation (Nikitas et al., 2025). Where in a metropolitan network such a facility should be placed is thus a question that has to be answered before its effect on cycling can be realized, and it is the question this thesis takes up. Section 2.2 characterizes the concept in full.

Two developments widen the setting in which such a facility is planned. The first is the electric bicycle, which extends the distance over which cycling competes with the car and brings longer commutes within reach (de Kruijf et al., 2018; Huang et al., 2024). The second is that support for cycling is increasingly organized at nodes rather than along links, in the form of mobility hubs where shared bicycles, shared cars and public transport meet (Nikitas et al., 2025). Such nodes make the bicycle part of a chain as well as a mode of its own: a ride to a station or transit hub followed by train or bus is a realistic alternative to the car for commutes too long to cycle in full (Tavassoli and Tamannaei, 2020). The Metropolitan Region Eindhoven (MRE) is moving in both directions, investing in regional hubs alongside cycling infrastructure and employer-oriented measures (Metropoolregio Eindhoven, nd), while Eindhoven works toward a citywide network in which a hub with shared transport lies within 200 to 300 meters of every address (Gemeente Eindhoven, nd). A CSP is the cycling-specific member of that family of nodes, and it faces the same siting question, namely where in a network a supporting node earns its place. The framework developed here is therefore built on commuter cycling but is not tied to the traditional home-to-work bicycle trip, and Section 2.3 sets out this positioning in full.

1.2 Problem Statement

Despite growing investment in cycling infrastructure, urban planners and infrastructure developers lack a systematic, demand-informed approach to deciding where cyclist support facilities should be placed. The case for such facilities is strongest along the route rather than only at its ends. Weather and topography are largely beyond planning control, but their deterrent effect can be reduced by sheltered rest points, changing facilities and charging for electric bicycles, whose motor assistance offsets the effort of hills and extends the viable cycling range (Parkin et al., 2008; de Kruijf et al., 2018). The CSP is designed around this set of en-route amenities (Loendersloot, 2022), which means that its value depends on being placed on the routes that cyclists actually ride.

The realization of such facilities is held back by several factors, among them the novelty

of the concept, the cost of construction and the question of who operates a facility once it is built. This thesis addresses one of them, the absence of an evidence-based method for deciding where the facilities should go. That difficulty is compounded by scarce and fragmented data on cyclist behavior at the metropolitan scale. Research on support facilities has concentrated on whether they encourage cycling, and much less on where in a network they should be placed (Chapter 2). Without a replicable framework that combines spatial analysis, behavioral simulation and demand modeling, decisions about where to place CSPs risk being ad hoc, inefficient or misaligned with the needs of cyclists.

1.2.1 Contribution

This research bridges transport modeling and urban planning with an academic and a societal contribution.

Academic Contribution

Academically, this study advances the methodology of active mobility planning by establishing a transparent, demand-informed framework that operates at the street level. While existing literature often focuses on macro-level mode choice or lane design, this thesis contributes an approach to generating synthetic, high-resolution cycling demand on a detailed street network. It couples this street-level demand, generated with an agent-based transport simulation, to a multi-criteria, flow-capturing optimization model. The result is a fully inspectable, rules-based method for siting supporting infrastructure such as Cycle Service Points, in which every siting recommendation traces back to explicit criteria and an auditable optimization rather than to an opaque model.

Societal Contribution

Societally, the research delivers two products. The first is a regional picture of cycling demand, hourly and at the level of individual street links, which shows planners, policymakers and municipal authorities where and when cyclists travel. It can inform network prioritization and safety work independently of where any facility is placed. The second is a ranked set of locations for Cycle Service Points together with the method that derives them, which moves facility siting away from ad-hoc decisions toward an evidence-based choice of where and how many to build. Together the two support the wider societal aims of lowering practical barriers to cycling to work, encouraging a shift away from car dependency, and helping cities accommodate population growth with sustainable, health-oriented transport.

To turn these aims into a workable study, the next section states the central research question and the objectives that structure the thesis.

1.3 Research Objectives

1.3.1 Central Research Question

> How can a demand-informed framework be developed to determine optimal Cycle Service Point locations, and where should Cycle Service Points be placed in the metropolitan region of Eindhoven to best support cycling commuters?

This thesis aims to investigate where to place Cycle Service Points based on a demand-informed approach, and to develop a framework that makes this determination

replicable across different urban contexts. To achieve this overarching aim, the following research objectives are defined:

1. To develop a generalizable demand-informed framework for CSP location planning, grounded in spatial analysis, behavioral simulation, and multi-criteria location optimization, that is applicable across urban regions in the Netherlands and provides a transferable basis for use beyond it.
2. To analyze the spatial characteristics of cycling networks, including lane types, land use, and connectivity, as inputs to the framework for identifying high-potential zones for CSP placement.
3. To simulate cyclist behavior and generate synthetic demand datasets using Multi-Agent Transport Simulation (MATSim), addressing the limited availability of empirical cycling data at the metropolitan scale.
4. To apply the framework as a case study in the metropolitan region of Eindhoven, producing concrete CSP location recommendations.
5. To evaluate the framework's outputs against real-world constraints, including stakeholder priorities, and reflect on the conditions under which it can be transferred to other geographical contexts.

1.4 Scope

This research operates within a defined scope structured across geographic, behavioral, and functional dimensions. Geographically, the overarching planning framework is designed for general application across urban regions in the Netherlands by leveraging uniform, national open-source spatial datasets. However, its practical execution is explicitly bounded to a single case study within the MRE. Consequently, the empirical findings and concrete facility optimization recommendations produced by this study are spatially constrained to this specific region.

Behaviorally, the target population is strictly limited to utility and commuter cyclists, with mostly home-to-work and home-to-education journeys. While a broader societal goal of implementing CSPs is to encourage a modal shift among car-dependent commuters, the optimization models evaluate travel patterns generated by utility travelers to ensure facilities intercept high-volume cycling corridors.

Methodologically, the study is bounded by the inherent sparsity of empirical active travel data at the regional scale. To establish continuous spatial coverage, synthetic datasets generated through transport simulation are utilized to model demand. While these models are grounded in transport behavior theory, the underlying routing logic relies on deterministic assumptions regarding cyclist preferences and network speeds that may not fully capture dynamic, real-world cyclist heterogeneity. Therefore, the resulting demand surfaces and optimization outputs must be interpreted as strategic planning-level guidance rather than exact empirical prescriptions.

Finally, the operational scope does not extend to the detailed architectural design, structural engineering, or financial cost modeling of individual CSP facilities.

1.5 Thesis Structure

The remainder of this thesis is organized as follows. Chapter 2 presents a review of the relevant literature, covering cycling behavior, cycling support facilities, mobility hubs, trip-generation regression, agent-based modeling and spatial optimization methods.

Chapter 3 defines the demand-informed framework, including the bicycle network, the MATSim demand simulation, and the flow-capturing siting model. Chapter 4 applies Phase 1 of the framework to the MRE, reporting the calibrated demand model and its validation, and Chapter 5 applies Phase 2, reporting the suitability surface and the recommended CSP locations. Finally, Chapter 6 answers the research question and sets out its practical and policy relevance, evaluates the model against the literature, and closes with the limitations of the work and directions for future research.

Chapter 2

Literature Study

This chapter reviews the literature that underpins the framework developed in this thesis, in six parts that follow the order in which the framework is built. Section 2.1 reviews what makes people cycle, and commute, by bicycle. Section 2.2 introduces the Cycle Service Point (CSP) concept and the support facilities it combines, and Section 2.3 places it within the wider literature on mobility hubs. Section 2.4 covers the regression models used to estimate how many cycling trips each zone produces, and Section 2.5 the agent-based simulation that turns those trips into routes on the network. Section 2.6 reviews the location models that decide where facilities go, and Section 2.7 draws these strands together and identifies the research gap.

2.1 Cycling Behavior and Mode Choice

The decision to cycle particularly for commuting is shaped by a complex and interacting set of individual, environmental, and contextual factors. Heinen et al. (2010) provide a comprehensive review of the cycling commuting literature, identifying five broad categories of determinants for commuting: the built environment, the natural environment, socio-economic characteristics, psychological factors, and travel cost, time, effort, and safety. Their review remains one of the most widely cited syntheses in this field and provides the theoretical backbone for understanding cycling mode choice.

2.1.1 The Built Environment: Distance, Density and Infrastructure

Trip distance is consistently identified as the most important built-environment determinant of cycling. As commuting distance increases, cycling's share of modal choice falls sharply. Heinen et al. (2010) reports that non-cyclists mention lengthy distances as a barrier (27% vs. 3% of cyclists), while cyclists tend to live closer to work. In the same article, it suggests that women cycle shorter distances to work than men, with reported thresholds of approximately 6.6 kilometers for women and 11.6 kilometers for men.

Urban form characteristics, particularly residential density, land-use mix, and network connectivity also play a meaningful role. Higher residential densities are generally associated with shorter trip distances and lower car ownership, both of which positively influence cycling. Mixed land uses reduce the need for long-distance travel, increasing cycling's share of mode choice (Cervero et al., 2009). The quality and continuity of cycling infrastructure matter significantly: cyclists consistently prefer dedicated cycle paths to on-road lanes or shared roads, and discontinuities in infrastructure are perceived negatively, particularly by inexperienced and female cyclists. Constructing bicycle paths has been found to increase the cycling share at some locations by 1–2%, though the effect is

context-dependent and may be smaller in areas where cycling is already prevalent (Heinen et al., 2010).

Facilities at the destination particularly secure bicycle parking and, to a lesser degree, showers and changing facilities are identified as important facilitators of cycling commuting. The strongest preference is for bicycle lockers, followed by enclosed storage and open racks. Although cyclists value these facilities, the empirical evidence on whether their presence actually increases cycling frequency is mixed, suggesting that endpoint facilities are necessary but not sufficient conditions for cycling growth (Heinen et al., 2010).

2.1.2 The Natural Environment: Weather, Topography and Season

Unlike motorized transport, cycling is acutely sensitive to weather and landscape. Precipitation is consistently identified as the most deterring weather factor, with rain mentioned as the most negative aspect of cycling by commuters in multiple countries (Heinen et al., 2010; El-Assi et al., 2017). Women, recreational cyclists, and those who cycle in winter show a particularly strong aversion to rain. Seasonal variation in cycling is substantial: cycling shares are highest in summer and autumn and fall considerably in winter months, with the magnitude of decline varying by regional climate. In Sweden, the maximum acceptable cycling distance was found to fall from 20 km in summer to 10 km in winter (Bergström and Magnusson, 2003), illustrating how weather directly constrains the viable commuting range.

Hilliness has a negative effect on cycling for most users. Cities with steep topography consistently show lower cycling mode shares than flatter counterparts, with one UK comparison showing cycling shares of 13.1% in a relatively flat city versus 0.8% in a predominantly hilly one (Parkin et al., 2008). These natural environment factors are largely beyond planning control, but they create a clear rationale for infrastructure and amenities such as sheltered rest points, changing facilities, and e-bike charging stations that can mitigate their deterrent effects: covered rest points shorten exposure to rain, and charging support sustains the electric bicycles whose motor assistance offsets the effort penalty of hills and extends the viable cycling range (de Kruijf et al., 2018). The Cycle Service Point concept is designed around precisely this set of en-route amenities (Loendersloot, 2022).

2.1.3 Socio-Economic and Psychological Factors

Socio-economic factors show complex and context-dependent relationships with cycling. Most research finds that men cycle more than women, though this gender gap narrows or reverses in countries with strong cycling culture, such as the Netherlands and Belgium (Heinen et al., 2010). The relationship between income and cycling is ambiguous: high income enables bicycle ownership but also car ownership, and the net effect depends heavily on local norms and transport system characteristics (Heinen et al., 2010). Car ownership is among the strongest negative predictors of cycling mode share, while bicycle ownership has a positive one (Heinen et al., 2010; Parkin et al., 2008). This is the empirical basis for including car ownership among the candidate predictors of the trip-generation model of Section 3.4.4.

Psychological factors, particularly attitudes, social norms and habits, have emerged as highly influential. Cyclists hold more positive attitudes toward cycling and perceive fewer barriers than non-cyclists, and positive social norms in a workplace or neighborhood, such as colleagues who also cycle to work, increase the likelihood that an individual does so (Heinen et al., 2010). Habit is the harder barrier: people accustomed to driving tend not to consider cycling even where it is objectively feasible, so breaking the habit through targeted interventions, infrastructure or incentives is a central lever for promoting cycle commuting (Heinen et al., 2010). None of these attitudinal variables is observable in the

travel-survey microdata this thesis uses, so they enter the model only through the revealed behavior of the respondents who already cycle.

2.1.4 Work-Related Factors and Employer Influence

Beyond individual attitudes and built-environment characteristics, the workplace itself constitutes a significant and often underappreciated determinant of cycling mode choice. Heinen et al. (2013) investigated this dimension through a large internet survey of over 4,000 commuters across four Dutch municipalities, using binary logit models to identify which work-related factors predict both the decision to cycle and the frequency of cycling.

Cycling decisions are shaped by a cluster of employer-provided conditions. Bicycle storage inside the building, access to changing and shower facilities, and a public transport stop within 500 m of the workplace all independently raise the likelihood of cycling to work, while free car parking, a free car or a free public transport pass all lower it (Heinen et al., 2013). The simultaneous provision of cycling benefits and a free public transport pass is counterproductive: free public transport reduces cycling frequency even where cycling facilities are also offered (Heinen et al., 2013).

Social norms within the workplace exert an independent influence. Employees whose colleagues expect them to commute by car are significantly less likely to cycle than those whose colleagues expect them to cycle, an effect operating through both direct expectation and the organizational culture embedded in facility provision and financial incentives (Heinen et al., 2013). Heinen et al. (2013) further separate the determinants of becoming a cyclist at all from those of cycling every day: working hours, shift patterns and dress code predict full-time cycling frequency, whereas facility provision matters more for crossing the initial threshold. The same threshold logic appears in the wider commuting literature, where endpoint facilities are consistently named among the conditions under which longer cycle commutes become viable (Heinen et al., 2010; Hull and O'Holleran, 2014).

These findings are directly relevant to the CSP concept, which functions partly as an employer-adjacent cycling support facility. If the presence of changing facilities and secure storage at the destination matters for cycling uptake, then en-route support facilities addressing equivalent barriers along the journey rather than only at endpoints represent a complementary intervention that could extend the impact of employer-side investment. The CSP's changing rooms, lockers, and e-bike charging address barriers that endpoint facilities alone cannot resolve for longer-distance commuters.

2.1.5 Incentives, Long-Distance Cycling and Facility Needs

Research increasingly demonstrates that incentives and supportive conditions can extend the appeal of cycling beyond traditionally short distances. Studies show that even for journeys exceeding 20 km, a substantial share of commuters choose cycling over car use when the right conditions (high-quality dedicated infrastructure, financial incentives, and endpoint facilities) are in place (Heinen et al., 2010; de Kruijf et al., 2018). de Kruijf et al. (2018) evaluate a Dutch incentive program for the shift from car commuting to e-cycling and find a substantial increase in cycling under payment, which is the clearest evidence that the longer-distance segment responds to supportive conditions rather than to distance alone. Employer-provided facilities such as showers, lockers and secure parking are consistently rated as important by cyclists, although their provision in isolation does not reliably raise cycling frequency (Heinen et al., 2010, 2013).

This finding is particularly relevant to the CSP concept. The evidence suggests that commuters are willing to travel longer distances by bicycle when the journey is well-supported, but that the absence of practical endpoint and en-route facilities remains a significant deterrent especially for new or infrequent cyclists. CSPs are designed precisely

to provide such a supportive environment at strategically located points along commuting corridors, reducing the practical and psychological barriers that limit the cycling catchment for longer-distance journeys.

2.2 Cycling Infrastructure and Commuter Support Facilities

Section 2.1 showed that support at the workplace and along the route affects whether commuters cycle, particularly over longer distances. This section describes the facility concept that bundles such support. To systematically address the infrastructural deficits faced by active commuters, the CSP concept establishes a formalized architectural and functional framework for bicycle-centric support nodes. Developed by Marcel van Beek at Amsterdam Airport Schiphol as part of the goinGDutch consortium, in collaboration with BAM Infra Nederland, the infrastructure is categorized into distinct operational layers depending on the local scale, spatial constraints, and regional demand dynamics (Loendersloot, 2022).

The functional matrix of a standard CSP is bifurcated into primary and secondary services:

- **Primary Functions:** Core operational amenities directly facilitating the commuter's trip continuity, including secure bicycle parking, self-service and professional repair stations, rapid e-bike charging infrastructure, and smart changing facilities equipped with secure lockers.
- **Secondary Functions:** Value-added services intended to maximize node utility and dwell-time productivity, ranging from parcel collection kiosks and basic catering to integrated co-working spaces or micro-event venues.

The spatial logic governing CSP deployment relies heavily on high-intensity cycling corridors and critical multimodal transfer nodes, as demonstrated by the Schiphol Noord pilot infrastructure (Loendersloot, 2022). That pilot is also the only realized instance at the time of writing, and it supplies the order of magnitude of the investment a single facility represents (BAM Infra Nederland, 2025). The concept is therefore evidenced by a single built example rather than by a body of operating facilities, which is why this thesis can validate its demand model against observed counts but cannot validate its siting model against realized locations.

The functions the CSP bundles map onto the barriers the commuting literature identifies. Secure parking, changing facilities and showers are precisely the endpoint provisions that Heinen et al. (2013) find raise the probability of cycling to work, and e-bike charging addresses the range constraint that governs the longer-distance segment the concept targets (de Kruijf et al., 2018). The difference is one of position rather than of function: a CSP offers them en route rather than at the destination, which is what makes where it stands a research question rather than a matter of building regulation.

Consequently, to build a robust location-allocation model, the CSP concept is best contextualized within the broader mobility hub literature. Merging these localized active travel principles with the spatial optimization techniques reviewed further in this chapter provides the toolkit required to execute demand-informed CSP siting.

2.3 Mobility Hubs and Multimodal Transport Integration

The CSP concept sits within the broader academic and policy discourse on mobility hubs: visible, accessible, and integration-enabling spaces where public transport, shared mobility, micro-mobility, and active travel modes co-exist alongside supporting infrastructure. The connection is structural rather than nominal. A mobility hub and a CSP share the same

organising logic, concentrating supporting infrastructure at accessible, well-located nodes in order to lower the practical friction of using a sustainable mode, and they share the same central siting question of where in a network such a node earns its place. A CSP can be read as a hub specialised to a single mode: where a general mobility hub brokers transfers between many modes, a CSP removes the barriers that interrupt a cycling trip already under way. Positioning the concept this way is what allows the mature hub literature, its evidence on siting, adoption, social value and the risks of getting placement wrong, to be carried over directly into the CSP problem. Nikitas et al. (2025) provide the most comprehensive recent synthesis of this literature, conducting a systematic review and identifying the key opportunities and risks associated with hub implementation through a thematic narrative analysis.

2.3.1 How a Cycle Service Point Differs from a Mobility Hub

The two concepts are easy to conflate, because a large CSP and a small mobility hub can look alike on a plan: both are compact, both cluster amenities, and both sit beside a cycle route. They differ in four ways that matter for siting.

The first is the function of the stop. A mobility hub exists so that a traveller can change mode, and its value is the set of connections it brokers (Nikitas et al., 2025). A CSP exists so that a cyclist can carry on, and its value is the set of interruptions it removes from a trip that stays on the bicycle throughout (Loendersloot, 2022). Parking, repair, charging and changing are all services to an uninterrupted cycling trip, not to a transfer.

The second follows from the first. Hub demand is access and egress demand: users travel to the hub because it is where the transfer is possible, so the relevant catchment is the population or the jobs around it. CSP demand is pass-by demand: users are already in motion on a route chosen for other reasons, and the facility is used only if it happens to lie on that route. This is the classical distinction between a facility people travel to and a facility people pass (Hodgson, 1990; Berman et al., 1992), and it is the reason the siting model in this thesis captures flow rather than covering residents.

The third is modal scope. A hub is multimodal by definition, and much of the hub literature is concerned with the interaction between modes, with shared vehicles and with the public-transport connection (Nikitas et al., 2025). A CSP serves one mode. Questions that dominate hub research, such as which shared modes to co-locate or how to price a transfer, do not arise.

The fourth is the revenue model. Hubs are typically public or public-private and are justified on network accessibility (Nikitas et al., 2025), whereas the CSP pilot is underwritten by employer subscriptions (Loendersloot, 2022), which makes the employment base near a position a siting criterion in its own right rather than a by-product of centrality.

Where the two concepts do overlap is in the empirical lessons: the hub literature's evidence on adoption, on the risk of a facility going unused and on the social value of a well-placed node transfers to the CSP case directly, and it is used that way in this thesis. Its siting methods do not transfer, and the remainder of this chapter builds the alternative.

2.3.2 Defining and Typologising Mobility Hubs

Mobility hubs have been defined in diverse ways across the literature, from simple physical interchange points connecting multiple transport modes to ambitious integrated urban nodes offering transport, commercial, social, and environmental functions (Nikitas et al., 2025). Definitions converge on the ideas of modal integration, multimodality, and accessibility as core functions, with the specific mix of modes and amenities varying by scale and context. The concept has antecedents in Dutch Park and Ride policy from the mid-1970s, but has grown substantially in prominence over the past decade due to increasing emphasis on sustainable transport planning and carbon emission reduction.

Contemporary examples include mobil.punkte in Germany, WienMobil in Austria, and Hoppinpunten in Belgium (Nikitas et al., 2025).

Hub typologies range from large transit-oriented development nodes in city centers, to neighborhood community hubs, to small ancillary single-mode stations on the urban periphery. In sub-central and suburban areas, mobility hubs can serve as first- and last-mile connectors to major public transport, piecing together a city's mobility network. In rural contexts, even minimal hub provision a bike-share or car-share point with demand-responsive transport links can meaningfully improve accessibility where public transport is sparse (Nikitas et al., 2025).

2.3.3 Opportunities: Multimodality, Accessibility and Social Value

The primary opportunity associated with mobility hubs is their capacity to enable multimodal and intermodal travel, providing a viable alternative to private car use. By integrating transport options from conventional public transit to shared micro-mobility and bicycle facilities, hubs reduce the cognitive and practical effort of multimodal travel (Nikitas et al., 2025). Nikitas et al. (2025) find evidence that frequent hub users become more likely to reduce car ownership and acquisition, particularly when hubs are embedded within a network and supported by real-time information provision and digital mobility services such as Mobility-as-a-Service (MaaS).

Mobility hubs also carry social functions beyond transport. Amenities such as green spaces, co-working facilities, commercial services and family-oriented infrastructure can make hubs genuine community assets, and architectural quality matters more to users than previously assumed: in a Zurich study, architectural identity was rated more highly than the provision of car sharing or micro-mobility, with design quality influencing both the appeal of the hub and attachment to it (Nikitas et al., 2025). This social abundance dimension is highly relevant for CSPs, which similarly aim to offer services that extend beyond mere bicycle parking.

2.3.4 Risks: Adoption Barriers, Environmental Limits and Inequity

Nikitas et al. (2025) identify five major risk categories: environmental burden, public indifference, discomfort by default, abuse of space, and inequity. These are not merely theoretical. Evidence from Vienna's WienMobil stations shows that even well-located hubs can go largely unnoticed and underutilized where habitual car use is entrenched and shared mobility is unfamiliar (Nikitas et al., 2025). The modal shift achievable by hubs has also been questioned: in Munich, 70% of transport substitutions replaced public transport rather than private car trips, and other studies observed personal car-use reductions of only 0.9% and 4.6% following hub introduction (Nikitas et al., 2025). The same substitution question is unresolved for cycling facilities generally, where provision is not reliably followed by a measured increase in use (Heinen et al., 2010; Hull and O'Holleran, 2014).

Transfer discomfort represents a persistent structural challenge. Even where waiting facilities are excellent, users reveal a preference for no transfer at all, so the convenience of private car travel sets a high benchmark (Nikitas et al., 2025; Sheller and Urry, 2000). Social equity is a further concern: hubs placed primarily near high-demand urban nodes risk deepening the infrastructure gap between city centers and lower-income or rural areas, while digital interfaces and MaaS subscription models may exclude older or less digitally literate users (Nikitas et al., 2025).

These findings carry direct implications for CSP planning. The risk of public indifference underscores the importance of demand-led siting: placing CSPs where cycling demand is demonstrably high rather than where installation is convenient. The equity concern points to the need for equitable network-level distribution. The environmental burden risk particularly relevant to e-bike charging highlights the importance of ensuring that CSPs

genuinely support modal shift rather than merely redistributing existing sustainable travel behavior.

2.3.5 Bicycle Sharing as a Feeder Mode: Hub Network Design

A body of emerging research treats the integration of bicycle sharing with public transport not merely as a planning aspiration but as a formal network design problem, seeking to optimize the placement of bicycle sharing infrastructure to maximize modal shift away from private car use. Tavassoli and Tamannaei (2020) developed a mathematical programming model for this purpose, proposing a hierarchical hub network structure for Bike-and-Ride (B&R) services in which bicycle sharing stations at two levels those co-located with public transport stops (level-I hubs) and those distributed in surrounding neighborhoods beyond walking distance of transit (level-II hubs) are linked by designated bicycle lanes under a budget constraint.

The objective function of the model is to maximize the number of private car trips that shift to B&R, making it explicitly competitive in character: the model sites infrastructure only where the combined travel cost of cycling to a hub, riding the transit network, and cycling to the destination is lower than the cost of driving (Tavassoli and Tamannaei, 2020). The model was applied to Isfahan, Iran, where 40% of daily trips are made by private car and transit catchment areas leave over half the population uncovered. With a budget of $30 million, the optimized B&R network was found to shift 38% of eligible private car trips, representing 16.7% of total private car trips, to the B&R mode, while extending the catchment area of rapid transit services by over 50% (Tavassoli and Tamannaei, 2020).

The sensitivity analysis conducted by Tavassoli and Tamannaei (2020) yielded several findings relevant to cycling infrastructure planning more broadly. Shifted demand grows with budget but at a diminishing rate, so beyond a certain network density improvements must be sought in fare or service quality rather than further expansion. Subsidies on transit fares produce greater modal shift than equivalent subsidies on bicycle sharing fees, because transit fares are paid more frequently per trip in a B&R journey. Passengers in higher-income societies are less responsive to monetary incentives where the car remains the faster mode, so B&R strategies work best where the combined mode is time-competitive with driving or where car travel times are deliberately raised through demand management (Tavassoli and Tamannaei, 2020). The diminishing-returns result recurs in this thesis, where the marginal flow captured by each additional CSP falls as the set grows, a property that follows from the objective rather than from the budget (Nemhauser et al., 1978).

For CSP planning, the hierarchical hub network framework offers a structural parallel. Like level-II B&R hubs, CSPs extend the effective cycling catchment of longer-distance trips beyond the immediate vicinity of transit nodes. The two models part company on the demand object: Tavassoli and Tamannaei (2020) maximize car trips converted, which yields a modal-shift figure directly but requires a calibrated mode-choice model, whereas the formulation used here maximizes intercepted cycling flow, which needs no such model and stops one step short of a modal-shift estimate.

2.4 Trip Generation Regression for Cycling Demand Scaling

A fundamental challenge in cycling infrastructure planning at the metropolitan scale is the limited availability of observed cycling trip data. National travel surveys such as Onderweg in Nederland (ODiN) sample only a small fraction of trips: after filtering to intra-regional weekday bicycle trips, the 2022 wave used here retains on the order of 0.9% of its regional records (Section 3.4.2). Simulating this thin sample directly in a transport simulation risks under-representing the true volume and spatial distribution of cycling demand.

Trip generation regression addresses this gap by using a statistical model calibrated against ODiN observations to estimate total cycling trips generated by each zone, producing scaled totals that underpin a synthetic population of realistic size. The classical four-step travel demand model has used regression for trip generation since the 1960s, relating observed trip rates to measurable zone-level predictors such as population, employment density, car ownership, and land use composition (de Dios Ortúzar and Willumsen, 2011).

Because trip counts are discrete, non-negative integers with a right-skewed distribution, ordinary least squares regression is statistically inappropriate. The Poisson regression model assumes that the variance of the dependent variable equals its mean; however, cycling trip counts typically exhibit overdispersion, where the variance substantially exceeds the mean due to unobserved heterogeneity across zones in factors such as local cycling culture, infrastructure quality, and employer facilities (Cameron and Trivedi, 2013). The negative binomial regression model addresses this by introducing a dispersion parameter estimated from the data, and is empirically validated for cycling demand modeling (Faghih-Imani and Eluru, 2016; El-Assi et al., 2017).

Which predictors to retain is a model-selection problem rather than a significance test, since several candidate predictors of cycling demand are strongly inter-correlated. The standard criterion is Akaike's, which trades fit against the number of parameters and so penalises a predictor that buys likelihood only by adding a degree of freedom (Akaike, 1974). A count model's fit is then reported as a pseudo R^2, which compares log-likelihoods rather than explained variance and is not on the scale of a linear R^2 (Cameron and Trivedi, 2013). The two studies above report the same pattern: built-environment covariates are significant but account for a modest share of the variation once exposure is controlled (Faghih-Imani and Eluru, 2016; El-Assi et al., 2017).

The predictors, the selection procedure and the estimation used in this thesis are set out in Section 3.4.4.

2.5 Agent-Based Modeling and MATSim for Transport Simulation

A trip-generation regression fixes how many cycling trips each zone produces, but not where those trips go or which routes they follow. Continuous cycling counts are available only at selected points, origin-destination data for cycling trips are rarely collected at the granularity needed for corridor-level facility siting, and large-scale Global Positioning System (GPS) tracking datasets, though growing in availability, remain restricted in geographic and temporal coverage. Agent-Based Modeling (ABM) and transport simulation address this gap by generating synthetic demand grounded in behavioral assumptions and spatial constraints (Horni et al., 2016).

2.5.1 Agent-Based Modeling in Transport Research

ABM simulates individual travelers according to specified decision rules and observes the system-level patterns that emerge from their interactions. The aggregate four-step model works with zonal matrices and loses the individual once the trips are summed (de Dios Ortúzar and Willumsen, 2011), whereas an agent-based model keeps each traveller as a distinct entity with a full daily plan, which preserves heterogeneity and captures the non-linear dynamics that arise when many individuals make interdependent decisions at once (Horni et al., 2016). The approach has been applied across transport research, from evacuation planning and land-use interaction to car-sharing uptake and autonomous-vehicle assessment (Horni et al., 2016).

For cycling specifically, ABM lets the analyst simulate how individual cyclists choose routes and respond to infrastructure (Ziemke et al., 2017). Given an agent population with

realistic socio-demographic and spatial characteristics and a route-choice model calibrated against empirical evidence, the simulation yields a plausible synthetic dataset of cycling trips across a metropolitan network, including trips that no existing data source observes directly. That property is what this thesis needs: a facility siting model requires demand on every corridor, not only on the corridors that happen to carry a counter.

2.5.2 Route Choice Modeling for Cyclists

Route choice is a well-established subfield of transport research, and its application to cycling presents specific challenges and opportunities relative to motorized transport. Cyclists face a richer trade-off space than car drivers: they weigh not only distance and travel time, but also infrastructure type, surface quality, gradient, traffic exposure, continuity of dedicated facilities, and subjective safety perceptions. Heinen et al. (2010) summarize the empirical evidence on route preferences: cyclists consistently prefer dedicated paths over on-road lanes, prefer routes without parking-adjacent lanes, avoid discontinuities in cycling facilities, and dislike frequent stop signs and traffic signals, though the magnitude of these preferences varies by experience level and trip purpose.

Revealed preference and stated preference methods have both been applied to understand cyclist route choice. Revealed preference studies using GPS data provide the most ecologically valid evidence, showing that cyclists often take longer or hillier routes to avoid high-traffic roads or missing infrastructure links (Heinen et al., 2010), routes that would appear suboptimal from a distance-minimizing perspective alone. Ziemke et al. (2017) operationalize exactly this trade-off inside an agent-based simulation, adding surface, infrastructure type, comfort and gradient as link-level dis-utilities alongside travel time, which is the specification used in this thesis. Understanding these preferences is important for CSP siting: a CSP placed along a route that cyclists avoid, even if it lies on the geometric shortest path, will not attract the intended users. Demand modeling for CSP placement must therefore account for realistic route choice behavior, precisely the behavior that the agent-based simulation platform described next encodes at the link level.

2.5.3 MATSim: The Multi-Agent Transport Simulation Platform

Multi-Agent Transport Simulation (MATSim) is an open-source, Java-based ABM platform designed for large-scale transport network simulation (Horni et al., 2016). It implements a co-evolutionary algorithm in which agents iteratively improve their daily activity-travel plans through a process of plan generation, execution on a queue-based traffic simulation, and scoring (Horni et al., 2016). Agents are assigned an initial set of plans based on their home location, activity schedule and available modes; plans are executed on the network and scored with the Charypar–Nagel utility function, which rewards time spent in activities and penalizes travel (Charypar and Nagel, 2005); a fraction of agents then replaces its current plan with a newly generated or previously tried alternative; and this repeats across many iterations until the system reaches a relaxed state in which no agent can improve its score by unilaterally changing plan (Horni et al., 2016).

MATSim has been applied to cycling simulation in several contexts (Ziemke et al., 2017). The platform supports mode-specific network assignments, allowing cycling trips to be routed along bicycle-specific infrastructure with link-level travel time and discomfort parameters. Bicycle routing in MATSim can account for infrastructure type, gradient, and traffic exposure, enabling realistic synthetic generation of cyclist origin- destination flows and route distributions. The output of a MATSim simulation includes agent-level trip records origin, destination, departure time, route, and arrival time that can be spatially aggregated to produce corridor-level demand estimates suitable for CSP siting analysis. Where count data are sparse, such a simulation yields a demand surface that covers every corridor of the network rather than only the counted points.

2.5.4 Limitations of Synthetic Demand Generation

Synthetic demand inherits the uncertainty of its inputs. Its quality depends on the accuracy of the input population and activity schedules, on the calibration of the route choice model and on the specification of the utility function (Horni et al., 2016). Assumptions about travel behavior, such as the maximum acceptable cycling distance, the sensitivity to gradient or the value of cycling-specific infrastructure, are uncertain and need not reflect the heterogeneity of real commuter behavior (Ziemke et al., 2017). Simulated output is therefore planning guidance rather than empirical prediction, and validation against whatever counts exist is what makes the surface defensible. Which statistic to validate against is itself contested for an agent-based model, since the thresholds in routine use were written for motorized volumes (Highways Agency, 1996), and measures designed for this setting have been proposed in response (Friedrich et al., 2019; Scherr et al., 2020).

2.6 Optimized Location Search for Cycling Infrastructure

Given a demand surface, whether empirical or synthetically generated, the problem of locating cycling support facilities such as CSPs can be formalized as a spatial optimization problem (Drezner and Hamacher, 2002). The relevant literature spans classical facility location theory, the flow-capturing models developed for facilities that serve demand en route (Hodgson, 1990), the metaheuristic methods used when a siting objective has no exploitable structure (Gen and Cheng, 1997), the multi-criteria weighting of siting factors, and the spatial analysis of cycling networks. The distinction that organises this section, and that ultimately determines which model class is appropriate for a CSP, is whether demand travels to a facility or past it.

2.6.1 Classical Facility Location Theory

The classical p-median and p-center facility location problems provide the theoretical foundation for demand-weighted infrastructure siting. The p-median problem seeks to locate p facilities to minimize the total demand-weighted distance between each demand point and its nearest facility (Drezner and Hamacher, 2002). The p-center problem minimizes the maximum distance between any demand point and its nearest facility, ensuring equitable coverage (Drezner and Hamacher, 2002). Both formulations are NP-hard for general networks, necessitating heuristic or metaheuristic solution approaches for realistic problem instances (Drezner and Hamacher, 2002; Gen and Cheng, 1997).

In the transport context, facility location models have been extended to account for the spatial structure of networks, the directionality of flows, and the capacity constraints of facilities. Covering models including the Location Set Covering Problem (LSCP) and the Maximal Covering Location Problem (MCLP) frame siting in terms of the proportion of demand covered within a specified access distance (Church and ReVelle, 1974), and they are the standard tool for facilities such as clinics, schools or depots that users deliberately travel to. A defining assumption of this covering family, however, is precisely that: demand is attached to a fixed residential origin and a resident is counted as served once a facility lies within their access distance. That assumption does not hold for a CSP. A CSP is a discretionary, pass-by facility, used by cyclists already in motion along a route chosen for reasons of their own; siting it well means intercepting the flow on the corridors people actually ride, not enclosing the neighborhoods they depart from. Modeling a CSP as a covering problem would reward proximity to origins over presence on the route, which is why the appropriate model class is the flow-capturing family discussed next rather than the covering family.

2.6.2 Flow-Capturing Location Models

The flow-capturing location model was introduced by Hodgson (1990) to address exactly the demand structure that defeats a covering model: demand that is not a set of fixed points but a set of flows moving along paths through a network. A facility is rewarded not for lying near where trips begin, but for lying on the paths that trips actually follow, so that it captures the flow passing it. Berman et al. (1992) developed the same idea specifically for discretionary service facilities, those a traveller patronises while already en route rather than making a dedicated trip, which is precisely the behavioral category a CSP belongs to. In the original formulations capture is binary: a facility captures a flow if and only if it lies exactly on that flow's path. Berman et al. (2003) relaxed this to a gradual covering decay, in which capture falls smoothly with distance from the path, so that a facility a short distance off a corridor still intercepts part of its flow. This graded, distance-decaying capture removes the "coverage cliff" of a hard access radius.

A property of the flow-capture objective is important for how it can be solved. When each flow is credited to the single best facility that captures it, the total captured flow is a weighted coverage function, which is monotone and submodular: adding a facility never reduces the captured total, and the marginal gain of a facility can only diminish as more are already sited. For such objectives the classical result of Nemhauser et al. (1978) guarantees that a simple greedy algorithm reaches at least a $(1 - 1/e) \approx 63\%$ fraction of the optimum under a cardinality budget, and the accelerated and lazy-greedy refinements of Minoux (1978) and Leskovec et al. (2007) make the computation tractable at scale while supplying a tight, per-iteration bound on the remaining optimality gap. A flow-capturing formulation therefore admits a solution with a certified near-optimality guarantee, without recourse to the metaheuristics that unstructured objectives require (below). In the conventional formulations the flows themselves are origin–destination volumes routed on shortest paths.

2.6.3 Genetic Algorithms and GIS for Optimal Location Search

A foundational contribution to the integration of evolutionary computation with spatial analysis is the work of Li and Yeh (2005), who demonstrated that Genetic Algorithms (GAs) can be combined with Geographic Information Systems to solve complex multi-site, multi-objective facility location problems that are intractable for conventional optimization methods. Their key conceptual contribution is the encoding of candidate solutions as chromosomes in which each gene pair represents the coordinates of one facility location. This encoding allows the GA to optimize the combined placement of several facilities simultaneously, accounting for their interactions such as overlapping catchments in a way that sequential greedy heuristics cannot, while the fitness function remains flexible enough to combine multiple planning objectives, for example demand coverage, transport cost and proximity to infrastructure, through a linear weighted evaluation of GIS-derived inputs.

This capacity matters most where the objective has no exploitable structure. For combinatorial location problems at the scale addressed here, with $n > 500$ candidate sites, exact solution methods become computationally intractable (Drezner and Hamacher, 2002), and metaheuristics such as GAs are an established way to obtain near-optimal solutions within acceptable computation time for this problem class (Gen and Cheng, 1997; Deb, 2001). Metaheuristics are the appropriate tool precisely when interactions between facilities produce a non-additively-separable objective that precludes any guarantee from a sequential heuristic. Where the objective is monotone submodular, as for flow capture, a greedy algorithm already carries a provable bound that a metaheuristic does not offer.

2.6.4 Multi-Criteria Weighting of Siting Factors

Whichever location model is used, a CSP is scored on several spatial criteria at once, employment nearby, transit access, land-use mix, network centrality, and the aggregation of these into a single suitability value requires a set of weights. In Geographic Information System (GIS)-based multi-criteria siting the weights are routinely the least defensible part of the analysis, and in the genetic-algorithm siting model of Li and Yeh (2005) they determine the location directly, because the weighted criterion sum is itself the objective. Direct rating, in which the analyst asserts a percentage per criterion, may not be internally consistent. Structured-elicitation methods such as the Analytic Hierarchy Process improve on this through pairwise comparison, but they require a practitioner panel and a comparison set whose internal consistency degrades as respondents tire, and they still present a single elicited weight vector as though it were authoritative. A separate family sidesteps judgement altogether and derives the weights from the criterion data: the CRITIC method of Diakoulaki et al. (1995) weights a criterion by its contrast and by how little it duplicates the others. Importantly, none of these methods tests whether the criteria are independent, which is what makes the overlap between them a property worth measuring in its own right. How this thesis sets the weights is described in Section 3.5.4.

2.6.5 Spatial Analysis of Cycling Networks

Spatial analysis of cycling network characteristics including lane types, land use, intersection density, and connectivity provides a complementary input to demand-based optimization. Network centrality measures (betweenness, closeness, eigenvector centrality) have been applied to identify nodes and links that are disproportionately important to network performance, and which may therefore represent high-priority locations for infrastructure investment (Barthelemy, 2016). Computing them at metropolitan scale from open data is itself now routine (Boeing, 2017). Beyond network performance, street centrality has independent empirical support as a predictor of where retail and service activity succeeds: Porta et al. (2009) and Sevtsuk (2014) show that betweenness and reach centrality correlate strongly with the observed density of shops and services in dense urban settings, which is the mechanism by which a central position raises the commercial viability of a CSP. Land use indicators including proximity to employment, residential density, educational institutions, and public transport stops inform the latent demand for cycling and the likely user base for CSP facilities.

For CSPs specifically, the combination of network centrality analysis with simulated cycling demand provides a richer basis for siting than either approach alone. Locations with high betweenness centrality on the cycling network are structurally important regardless of current demand; locations with high simulated demand reflect the behavioral preferences of simulated commuters; and locations that score highly on both dimensions are the most compelling candidates for CSP development.

2.7 Summary and Identified Research Gap

The literature reviewed in this chapter establishes a clear foundation for the research approach developed in this thesis. Several key conclusions emerge.

First, the decision to cycle is shaped by a multi-layered set of determinants, with trip distance, infrastructure quality, weather conditions, and psychological factors playing the most significant roles. Heinen et al. (2013) extend this picture to the workplace context, demonstrating that employer-provided facilities, financial incentives, and organizational social norms independently predict both the decision to become a commuter cyclist and the frequency of cycling. These determinants inform not only where cyclists travel, but where they are most likely to need supportive facilities along their route. The evidence

suggests that practical endpoint and en-route facilities can extend cycling catchment areas and attract commuters who would otherwise choose car travel but only when placed in locations where cycling demand is already high or realistically latent.

Second, the CSP concept responds directly to the documented barriers to cycling commuting by providing repair, rest, and charging infrastructure for e-bikes at strategically located nodes. It shares conceptual ground with mobility hubs but is specific to cycling commuters and oriented around en-route support rather than multi modal interchange. The mobility hub literature provides both inspiration and caution: the risks of public indifference and modal appropriation documented by Nikitas et al. (2025) underscore the importance of evidence-based, demand-led siting rather than supply-driven installation. The B&R hub network design of Tavassoli and Tamannaei (2020) shows that optimizing where cycling facilities are placed can reduce private car trips, even in car-dependent contexts.

Third, the demand-generation and location-optimization methods reviewed, negative binomial trip generation and MATSim for synthetic demand and facility-location theory for infrastructure siting, collectively provide the technical toolkit needed to execute CSP location planning at the metropolitan scale. The decisive distinction is between covering models, which attach demand to fixed origins, and flow-capturing models (Hodgson, 1990; Berman et al., 1992), which reward a facility for lying on the paths that trips actually follow. Because a CSP is a discretionary, pass-by facility, the flow-capturing family is the appropriate model class, and its objective is monotone submodular, so a greedy algorithm attains a certified near-optimal solution (Nemhauser et al., 1978). The genetic-algorithm and metaheuristic literature (Li and Yeh, 2005) applies to unstructured objectives that lack this property.

Fourth, and most importantly, no existing study has applied a systematic, demand-informed framework to the spatial optimization of CSPs specifically. The academic literature on cycling infrastructure focuses predominantly on lane design, mode choice determinants, and safety, with limited attention to the spatial placement of en-route support facilities along commuting corridors, and flow-capturing models typically take their flows from routed origin–destination volumes rather than from the link-level trajectories of individually simulated agents. This gap motivates the framework developed in Chapter 3, which integrates spatial cycling network analysis, MATSim-based demand simulation, and flow-capturing location optimization to produce the first replicable, evidence-based methodology for CSP siting applied as a case study to the Metropolitan Region Eindhoven.

Chapter 3

Methodology

3.1 Introduction

This chapter presents the methodology developed to address the central research question: How can a demand-informed framework be developed to determine optimal Cycle Service Point locations, and where should Cycle Service Points be placed in the Metropolitan Region of Eindhoven to best support cycling commuters? The study adopts a sequential two-phase quantitative design. The first phase models demand: it combines national travel-survey microdata with a zone-level regression model, a spatial destination-choice model and an agent-based traffic simulation in order to produce a spatially complete, link-level cycling demand surface for the study area. The second phase optimizes the siting of Cycle Service Points (CSPs) by scoring candidate positions on multi-criteria quality and selecting the set that intercepts the most quality-weighted cycling flow, a flow-capturing location problem solved on the simulated trajectories.

The Metropolitan Region Eindhoven (MRE), depicted in Figure 3.1, serves as the empirical context. Comprising the city of Eindhoven and its surrounding municipalities in the southern Netherlands, the region covers 1,457.7 km^2, features a dense cycling network and a high bicycle modal split, and therefore provides a suitable testbed for behaviorally grounded infrastructure optimization.



Figure 3.1: The MRE (study area).

3.2 Framework Overview

3.2.1 Aim of the Model

The model decides where, along a regional cycle network, a small number of service points should stand so that they lie on the routes the most cyclists already ride. Three properties follow from that aim and bound what the results can be used for.

It is a strategic instrument, not a forecasting one. The quantity it optimizes, intercepted flow, ranks candidate positions against one another. It does not predict how many cyclists stop, how many new cyclists a facility attracts or how many car trips it displaces, and none of those quantities is estimated anywhere in this thesis.

It models revealed demand under the current network. Routes are taken as given, so a position is credited only for cyclists whose existing path already passes it. A cyclist who would divert to reach a facility is not counted, which makes the coverage figures a floor rather than a central estimate.

Its output is a location at the resolution the demand supports. Neighbouring candidates on the same corridor capture almost the same flow, so the defensible recommendation is a stretch of corridor with a stated selection probability, and the choice of parcel within it belongs to a local feasibility assessment this model does not perform.

3.2.2 Structure of the Framework

The demand-informed framework deliberately separates spatial demand estimation from facility location planning. By focusing siting decisions on agent-level travel patterns, the optimization process reflects continuous spatial movement rather than static population aggregates. This also allows for independent verification of the demand model against observed bicycle counts before making any siting decisions. Figure 3.2 maps the framework on one page. It shows the Phase 1 and Phase 2 method blocks, the kind of result each block produces in Chapters 4 and 5, how the blocks feed one another, and the research objective each block answers.

Phase 1 (Section 3.4) is organized as a classical four-step travel demand model, re-specified for a single mode. Network Construction (Section 3.4.3) builds a bicycle-only routable graph from OpenStreetMap (OSM). Trip generation (Section 3.4.4) estimates how many cycling trips each zone produces and expands the thin survey sample into a synthetic agent population. Trip distribution (Section 3.4.6) assigns a concrete destination facility to every trip. Trip assignment (Section 3.4.7) routes the resulting itineraries over the network in Multi-Agent Transport Simulation (MATSim) until the system reaches a converged state, and validates the outcome against bicycle counters.

Phase 2 (Section 3.5) translates the simulated flows into concrete infrastructure recommendations by scoring candidate CSP positions on eight multi-criteria spatial factors and selecting the set that captures the most quality-weighted cycling flow, a flow-capturing location problem solved by a greedy algorithm with a certified optimality gap.

The methodology is modular by design. Every input is a nationally available open-data source and every processing step is implemented in open-source software, so the framework can in principle be reapplied to any other Dutch urban region by substituting the region-specific spatial inputs. This chapter defines that framework; its application to the MRE is reported as a case study in Chapters 4 and 5, and its transferability is revisited in the discussion (Chapter 6).

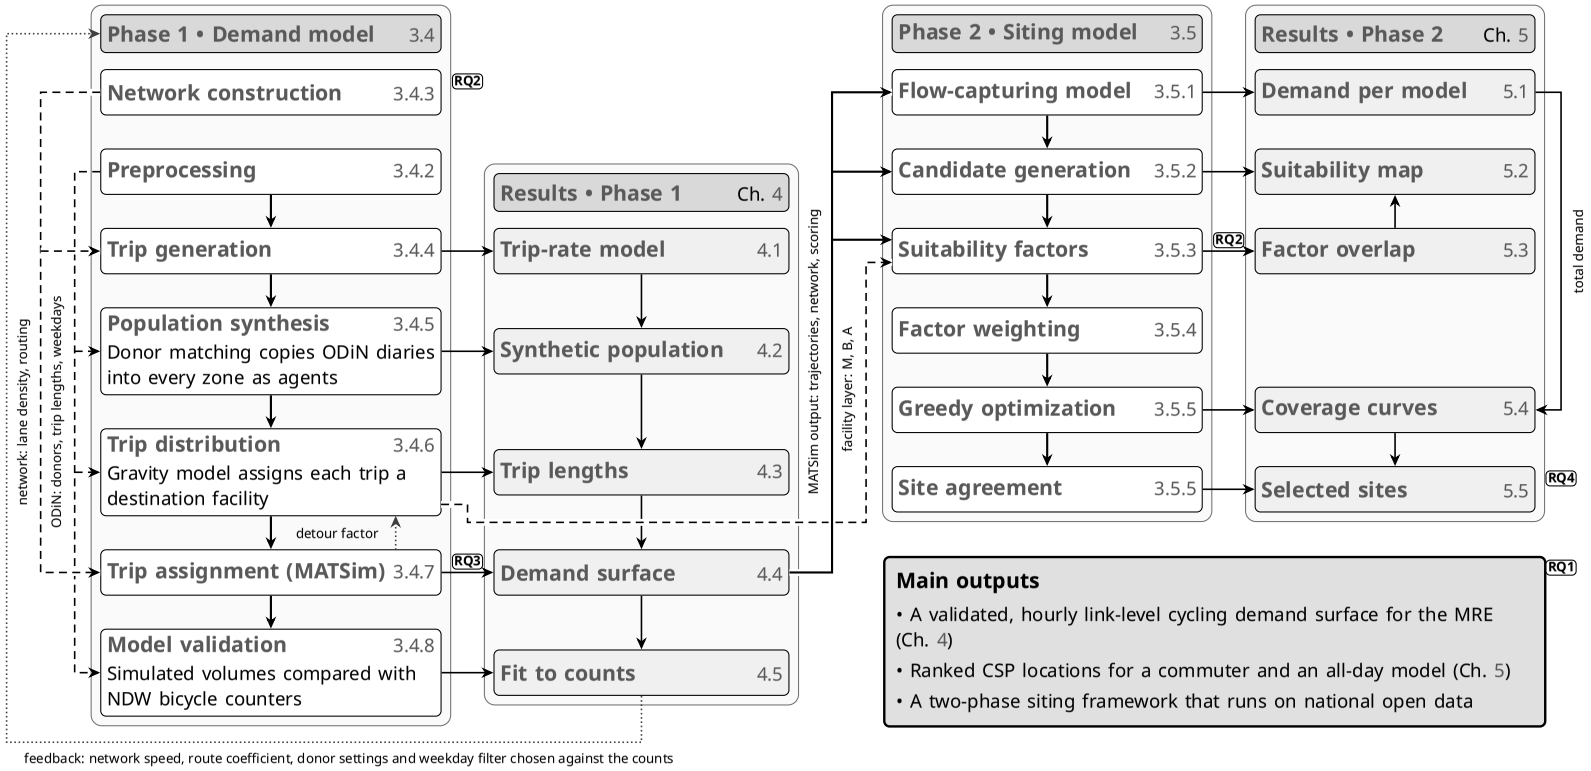

Figure 3.2: Framework of the thesis. Box headers link to their sections. Thick arrows give the order of steps and the hand-off of the demand surface to Phase 2. Thin arrows lead from a method to its results and between results. Dashed arrows are inputs a block passes to a later block, and dotted arrows are feedback, where a result was used to set an earlier setting. Labels name what is passed. RQn marks the block that answers research objective n (Section 1.3). RQ5 is answered in Chapter 6.

3.3 Data Sources

The data pipeline combines individual travel behavior, demographic indicators, land-use attributes and physical transport infrastructure from several open-data repositories.

Onderweg in Nederland (ODiN). The national travel survey conducted annually by Statistics Netherlands (CBS) on behalf of Rijkswaterstaat. The 2022 dataset contains about 62,000 respondents, with about 200,000 trips made. Each respondent reports a single-day trip diary that records, for every trip, the travel mode, trip purpose, origin and destination zones, departure and arrival times and activity duration, together with person- and trip-level expansion weights (`FactorP`, `FactorV`) that gross the sample up to the resident population. It covers the whole of the Netherlands and is released as anonymized microdata in successive annual waves (Centraal Bureau voor de Statistiek, 2023).

CBS spatial statistics. The official statistical geographies of the Netherlands, published by CBS as boundary geometries at the neighborhood (Buurt), district (Wijk) and four-digit postcode (4-digit postcode zone (PC4)) levels, each carrying demographic and socio-economic attributes such as population, household composition, income, car ownership and employment. The layers are released annually (Centraal Bureau voor de Statistiek, 2022b,a).

OSM. A collaboratively edited, openly licensed geospatial database of the world. It provides street- and path-network geometry together with descriptive tags for road classification, surface and cycling infrastructure, as well as points of interest. It is maintained continuously by contributors (OpenStreetMap contributors, 2024).

LISA employment register and the Bestand Bodemgebruik (BBG) land-use layer from CBS. Localized employment totals per sector and land-use composition metrics, used as zone-level regression predictors in trip generation and as spatial indicators in Phase 2 (Section 3.5) (Stichting LISA, 2023; Centraal Bureau voor de Statistiek, 2017).

Basisregistratie Adressen en Gebouwen (BAG) building register, via Publieke Dienstverlening Op de Kaart (PDOK). The official Dutch key register of addresses and buildings, maintained by municipalities under kadaster coordination. It records every addressable unit and building with its geometry, registered use class, official floor area and construction year, and is distributed as open data through the PDOK web feature service (Kadaster, 2024).

Nationaal Dataportaal Wegverkeer (NDW) bicycle counts. The NDW is the national road-traffic data portal, a partnership of Dutch road authorities that collects and openly publishes traffic measurements. Among these are permanent bicycle counting stations, which report bidirectional volumes at regular time intervals (Nationaal Dataportaal Wegverkeer, 2024).

3.4 Phase 1: Agent-Based Cycling Demand Model

Phase 1 produces the demand surface on which all of Phase 2 rests. It follows the structure of the classical four-step model, but every step is re-specified for a single mode and for disaggregate, agent-level output rather than zonal matrices. Mode choice, the third step of the conventional model, is absent by construction: the population consists exclusively of observed cyclists and the question is not whether they cycle but where and along which route.

1. **Pre-processing** (Section 3.4.2) processes the datasets by filtering, validating and interpolating the data into workable datasets.

2. **Network construction** (Section 3.4.3) builds the routable bicycle-only graph and the link attributes that drive route choice.

3. **Trip generation** (Section 3.4.4) estimates how cycling demand is distributed across zones, and calibrates its level to the survey control total.

4. **Population synthesis** (Section 3.4.5) expands the survey sample into a full population of individual agents, each carrying a complete daily activity chain.

5. **Trip distribution** (Section 3.4.6) converts each trip's zone-level destination into a concrete coordinate by means of a gravity-based destination choice model.

6. **Trip assignment** (Section 3.4.7) routes the resulting itineraries in MATSim and iterates until the system reaches a stable state.

3.4.1 Why MATSim Is Used for Assignment Only

MATSim can carry more of the four-step chain than it does here. Its replanning modules mutate an agent's mode, destination and departure time as well as its route, so one run could subsume trip distribution and mode choice (Horni et al., 2016). This model uses it for route assignment and departure-time mutation only, and generates demand upstream with a regression, a donor-based synthesis and a gravity model. Three reasons motivate that division.

What is being estimated. MATSim's choice modules are utility-driven: a destination is selected because it scores well under the Charypar–Nagel specification (Charypar and Nagel, 2005), so the parameters of that utility function become the object to calibrate. Calibrating destination choice inside the simulation would mean fitting those parameters against the same thin survey sample the model already uses, with no independent observation to fit them to. Estimating a gravity model on the survey's own observed trip-length distribution keeps the calibration target explicit and external to the simulation, and it is the standard treatment of the distribution step (Wilson, 1967; de Dios Ortúzar and Willumsen, 2011).

Mode choice is not a question this model asks. The synthetic population is drawn exclusively from respondents who made a bicycle trip, so every agent cycles by construction. Enabling mode choice would require a multimodal network, a car and transit supply model and a calibrated mode utility for each, none of which the CSP siting question needs and all of which would add error to the one output that matters here.

Method matched to purpose. Phase 2 requires a realized path per trip, resolved to the link. That is what the network-loading part of MATSim produces and what no aggregate method produces, since a four-step assignment returns link volumes without knowing whether two links carried the same cyclist. Route choice is also the one step where the agent-based treatment changes the answer rather than only the representation: identical origin–destination pairs take different paths in response to link attractiveness and departure time, which is the demand property Phase 2 exploits (Ziemke et al., 2017).

The cost is that routes are re-optimized inside the simulation while destinations are not, so the model cannot represent a cyclist who changes destination in response to the network, and the future-demand scenarios of Section 3.5.6 reweight existing trips rather than re-simulating travel choices.

3.4.2 Preprocessing

Spatial Framework

The spatial unit of analysis is the PC4 zone. The regional perimeter was established by intersecting the CBS PC4 polygons with the administrative boundary of the MRE which accounts to a total of 139 PC4 zones. All geometry is handled in the Dutch national coordinate reference system (RD New, EPSG:28992) so that every distance computation is metric.

The choice of PC4 as the spatial unit follows from the ODiN dataset, which locates each respondent's origin and destination at the PC4 level.

Travel-survey preparation

The ODiN preparation is deliberately split into two stages: a broad respondent base that is retained in full, and mode and geography filters that are applied downstream at the point of use. Figure 3.3 shows the flow: identification of the raw 2022 data, screening to the regional base, eligibility against the trip-level filters, and inclusion in the estimation sample.

The first stage keeps every respondent resident in Coördinatie Commissie Regionaal OnderzoeksProgramma (COROP) region 36 (Southeast Brabant), regardless of the modes they used, giving 7,295 diary records from 2,288 respondents out of the 200,054 records in the raw 2022 data. Filtering to bicycle trips at this point would be a mistake, even though every trip that is eventually simulated is a bicycle trip. The population synthesis of Section 3.4.5 weights candidate donors by how typical their household and demographic attributes are of a given zone, which requires the distribution of those attributes in the general population as a denominator. Restricting the base to cyclists makes that denominator unobservable: of the 2,288 respondents in the regional base, only 732 (32.0%) made at least one bicycle trip, so a cyclist-only base would discard 1,556 respondents' worth of donor-pool composition before any downstream step could see it.

The eligibility stage applies four criteria where they are actually needed, rather than up front. The principal travel mode (`KHvm`) must be bicycle (`KHvm` $= 5$). Both trip ends must lie within the study region (`VertCorop` and `AankCorop` in COROP 36), and the respondent's home postcode (`WoPC`) must fall inside the MRE perimeter. Records are indexed by respondent, so that the complete daily activity chain of every retained cyclist is kept rather than its bicycle legs in isolation.

Requiring both trip ends inside the region is a deliberately conservative choice. Applied to the 2,288-respondent regional base, the mode, geography and continuity filters together retain 1,852 intra-MRE bicycle trips, 0.93% of the raw 2022 data, which is the thin-sample problem that Sections 3.4.4 and 3.4.5 exist to address.

The model represents an average Monday–Friday day throughout. Diaries are filtered to `Weekdag` 2-6, taking the 2022 intra-MRE bicycle sample from 1,852 to 1,476 trip records (79.7%), and the counting-station observations used for validation (Section 3.4.8) are filtered to the same weekdays. Weekend cycling differs from weekday cycling in both volume and temporal profile, and is out of scope here because CSP siting is aimed at commuting demand. This Monday–Friday, 2022-only sample is the one taken forward into trip generation (Section 3.4.4).

Survey Weights and the Regional Control Total

ODiN attaches two grossing-up weights to every record: a person weight `FactorP` and a trip weight `FactorV`. The calibration uses both, and it uses them for different jobs, so it is worth saying first why the person weight is needed at all.

A travel survey is not a simple random sample of the population. CBS draws it with unequal selection probabilities by region and by household type, and response is not

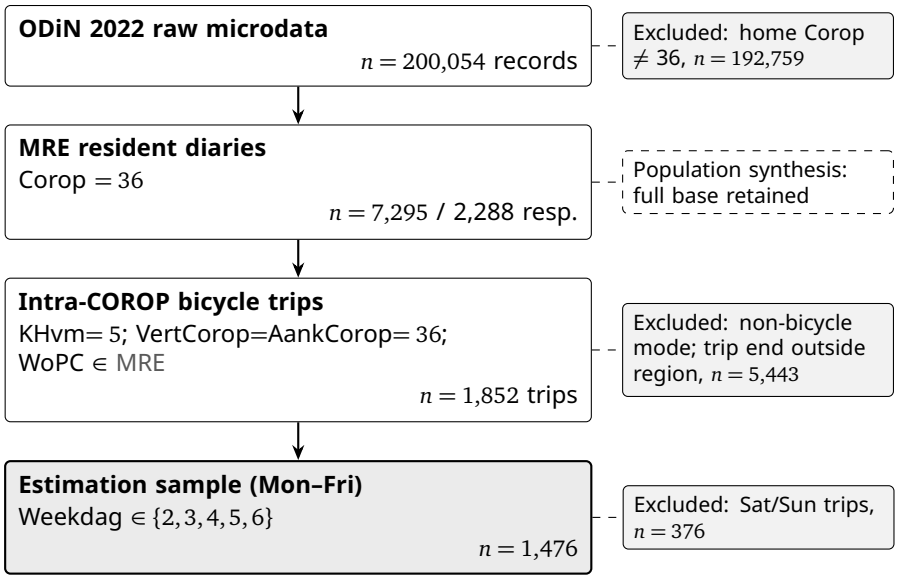

Figure 3.3: Construction of the 2022 Monday–Friday estimation sample from the raw ODiN microdata.

uniform across ages or areas, so a respondent in an under-sampled group stands for more residents than one in an over-sampled group (Centraal Bureau voor de Statistiek, 2023). FactorP is the factor that corrects for exactly this. Ignoring it and treating respondents as interchangeable would let the sampling design, rather than the region, determine which zones appear to cycle most, and since the trip-generation model of Section 3.4.4 is estimated zone by zone on a handful of respondents each, that distortion would propagate straight into the spatial pattern the whole model exists to produce. The person weight is therefore used in three places: to form the zonal cycling totals the regression is fitted to, to size the synthetic population, and as the quantity the borrowing procedure of Section 3.4.5 transfers to zones without a usable observation.

Two properties of FactorP make it the right weight for these jobs rather than FactorV. It is a person-level quantity, so it aggregates to a population that can be checked against an independent source, which is done below. And it is defined on a single assigned travel day, so a sum of person weights over a filtered set of days is still interpretable as a number of residents, whereas the trip weight is annual and has to be divided down before it means anything daily.

Each respondent r reports travel for a single assigned day and carries a person weight $\texttt{FactorP}_r$, the number of residents that respondent stands for. Summed over the complete seven-day sample, these weights reproduce the resident population. The model, however, is weekday-based, so the sample is first restricted to respondents whose travel day falls on Monday–Friday. Because respondents are spread roughly evenly across the week, this subset retains only about five-sevenths of the person weight and therefore sums to about five-sevenths of the population, not the full population. Multiplying by $7/5$ restores the full resident-population level, which is the quantity a weekday total must be compared against:

$$\hat{P}^{\text{wd}} = \frac{7}{5} \sum_{r \in \text{Mon–Fri}} \texttt{FactorP}_r, \tag{3.1}$$

where the sum runs over the Monday–Friday respondents and $\hat{P}^{\text{wd}}$ is the resulting population estimate.

Each trip v carries a trip weight $\texttt{FactorV}_v$, an annual weight: it grosses the recorded trip up to the number of equivalent trips made over the whole year. Restricting the sample

to Monday–Friday trips and summing these weights therefore gives the annual number of weekday trips; dividing by the 261 weekdays in the year converts it to an average-weekday total:

$$\hat{T}^{\text{wd}} = \frac{1}{261} \sum_{v \in \text{Mon–Fri}} \texttt{FactorV}_v, \tag{3.2}$$

where $\hat{T}^{\text{wd}}$ is the average-weekday trip control total.

These two totals (3.1,3.2) are the empirical anchors that tie the model to observed ODiN demand, and they enter the model at two different points. $\hat{T}^{\text{wd}}$ fixes the overall volume of travel to be reproduced: Section 3.4.4 calibrates the trip-generation model so that its predicted trip count matches this control total. $\hat{P}^{\text{wd}}$ fixes the number of agents that generate that travel: Section 3.4.5 sizes the synthetic population to it. Anchoring both to the same weekday basis keeps trips and travelers mutually consistent.

The two weights are internally consistent: on identical records the annual trip weight equals 365 times the person weight, so they describe the same population on different bases (a whole calendar year of trips versus an average day of persons). This is also why the control total, and not the raw regression fit, must set the demand level. The regression is fitted to a `FactorP`-based target, which is on a five-sevenths-of-population weekday basis, whereas the control total is a genuine weekday quantity; anchoring to the control total (Section 3.4.4) absorbs this basis difference automatically.

Before use, the person weights were checked against CBS population totals. Summed over the MRE respondents, `FactorP` reproduces the resident population aged 15 and over to within about 4% (a ratio of 1.04). The amount of respondents is near-exact in the youngest and oldest age bands and the deviation sits entirely in the 25–65 range. The model does not re-weight to remove this residual; because the demand level is fixed by the trip control total.

3.4.3 Network Construction

Route choice, travel time and every link-level output of the simulation are determined by the network the agents travel on. Because the model is bicycle only, that network cannot simply be a road network with a bicycle mode attached. Links (roads) a cyclist cannot or may not use have to be removed, and links a cyclist experiences differently from a car driver have to carry attributes the router can act on.

Clipping and Extraction

The network is built from a raw OSM extract of the Noord-Brabant region, converted to MATSim format and clipped to the MRE administrative boundary. The clipped, bicycle-accessible network contains 296,729 links. The bicycle-access filter removed 1,037,741 links from the extract. The tag vocabulary used to decide accessibility is derived from the CyclOSM rendering rules (CyclOSM contributors, 2024), which is the most widely used open codification of European bicycle access tagging. Table 3.1 lists the largest exclusion classes.

To make the network less detailed so runtime is optimized, node simplification is enabled. The link length is kept as the true length. This accumulates distance over every original OSM node and only resets when a segment is appointed. By simplifying the topology it reduces the number of links without shortening them, which would otherwise bias every travel time calculated. The original geometry is preserved as an attribute on 46.2% of links, so that simulated flows can be plotted on the real geometry rather than on node-to-node straight lines.

Links removed	Reason
542,873	`highway=footway` without explicit bicycle access
146,367	`bicycle=use_sidepath`
138,078	`bicycle=no`
109,062	`access=private`
43,337	`highway=pedestrian` without explicit bicycle access
30,464	`access=no`

Table 3.1: Largest exclusion classes applied by the bicycle-access filter, out of 1,037,741 removed links.

Link Attributes and Cycling Attractiveness

To govern bicycle routing behavior, network links are enhanced with three attributes derived from OSM data: physical comfort properties, facility classifications, and baseline free-flow speeds. Tags detailing surface type, smoothness, and bicycle specific infrastructure such as cycleways, cycle streets, and segregated paths provide the baseline data required by the bicycle contribution framework (Ziemke et al., 2017), with surface properties explicitly recorded across 47.8% of the network. These attributes are used because facility classifications directly inform the attractiveness scoring system.

Facility preferences are instead expressed through a link-level cycling attractiveness factor $f_a \in \{0.0, 0.5, 1.0\}$, which classifies each link as an official cycling route ($f_a = 1.0$), an ordinary cyclable link ($f_a = 0.5$), or a shared carriageway with fast motor traffic ($f_a = 0.0$). The routing engine translates this factor into a distance proportional dis-utility penalty,

$$C_a^{\text{attr}} = c_{\text{marg}} (1 - f_a) L_a, \tag{3.3}$$

where L_a represents the link length and c_{marg} denotes the marginal cost coefficient. The scale is capped at $f_a = 1.0$ to ensure that attractive routes minimize added disutility without introducing negative edge weights, which would result in an ill-posed shortest-path problem. Table 3.2 outlines the resulting distribution of these attractiveness scores across the network.

f_a	Link Class	Links
1.0	Official cycling route (national, regional, or local)	29,347
0.5	Ordinary cyclable link	266,840
0.0	Shared carriageway with fast motor traffic	542

Table 3.2: Distribution of the link attractiveness factor f_a across the 296,729 network links.

The final attribute category, free-flow speed, is held constant to prevent speed variations from distorting route selection. Although empirical speed profiles across facility types can be drawn from literature such as Schuhmacher et al. (2025), applying variable speeds introduces structural misalignments in the Dutch context, where for example, rural farm paths inadvertently inherit higher speeds than dedicated cycleways. To prevent speed differentials from penalizing high-quality cycling infrastructure, a uniform cruising speed of 21.1 km/h corresponding to the length-weighted mean speed of the network is assigned to all standard ridable links. Speed differentiation is retained solely where physical constraints dictate lower thresholds: service roads and living streets (17 km/h), footways, pedestrian areas and bridleways (15 km/h), and steps (3 km/h). Neutralizing free-flow speed in this manner ensures that route choices are driven mainly by network distance and the attractiveness formulation in Equation (3.3), with speed differences acting only on these slower link types.

Network Limitations

To establish a transparent foundation for interpreting flow assignments, three key structural characteristics of the network representation warrant explicit definition. First, directional routing operates strictly along the primary link topology. Although 2,553 links contain attributes indicating bicycle contraflow permissions (such as opposite-direction cycleway designations), these properties are recorded as descriptive metadata rather than converted into reversed graph edges. Route assignment therefore follows a single-direction topology, defining a clear boundary for directional network flows.

Second, the network applies a permissive inclusion rule for general path classifications to maintain comprehensive spatial coverage. A total of 26,508 links (representing 8.9% of the network) categorized as general paths are retained without requiring explicit bicycle tags. In the Dutch cycling context, this tag encompasses both paved dedicated paths and unpaved tracks; adopting this inclusive threshold ensures that genuine, high-utility cycling links are fully represented across the network rather than prematurely excluded.

Finally, the model configuration holds the comfort and surface quality coefficients at zero. Unpaved and soft-surface links therefore stay in the network and are routed in the same way as paved links of the same type.

3.4.4 Trip Generation

Trip generation answers one question: how is cycling demand distributed across the 139 PC4 zones of the study area? The survey observes a few hundred cyclists (Section 3.4.2). The simulation needs a full population across every zone. The obvious approach, scaling each zone up by its own sample, fails precisely where it matters. A zone with two respondents would be treated as a low cycling zone rather than as a thinly sampled one, and 26 zones with no respondents at all could not be scaled up by any factor. This method bypasses regional data scarcity by establishing a unified spatial relationship applied across all zone units through four sequential steps:

1. Standardize each zone's characteristics as deviations from the regional average to establish a common scale.

2. Estimate a single region-wide relationship between standardized attributes and daily cycling trips using only observed zones.

3. Apply the fitted relationship to all 139 zones, incorporating the 26 unobserved zones via register data.

4. Uniformly rescale zonal trip counts so the regional aggregate matches the survey weight control total.

This four-step workflow returns predicted daily cycling trips per zone, from which a person-level expansion factor is derived to determine agent generation in the population synthesis framework (Section 3.4.5). Essentially, the design separates demand magnitude from spatial distribution: total volume is anchored by survey weights, while geographic distribution is governed by the regression model.

Model Specification and Data Sample

The dependent variable, Y_i, represents the weighted number of daily weekday bicycle trips generated by residents of zone i ($i = 1, ..., 139$). Trips are attributed to the traveler's zone of residence rather than trip origin to avoid double counting individuals departing from multiple zones. Each zone yields exactly one aggregate observation.

Seven candidate attributes j ($j = 1, ..., 7$) cover categories established in the cycling demand literature (Cervero et al., 2009; El-Assi et al., 2017): population density, employment density, median income, car ownership rate, land-use mix, cycling lane density, and the share of residents aged 18–65. All candidate variables are standardized prior to estimation so that coefficients reflect the marginal effect of a one-standard-deviation change. Variable selection results are detailed in Section 4.1. The estimation sample comprises $n = 113$ zones with direct ODiN observations. The remaining 26 unobserved zones contribute zero weight to model optimization and are predicted using their registry attributes.

Regression Framework

Because cycling trips are non-negative count data, classical Ordinary Least Squares (OLS) regression is methodologically unsuitable. Standard Poisson regression is similarly restrictive due to its assumption of equidispersion (mean equals variance). Zone counts violate this assumption significantly, exhibiting substantial overdispersion driven by unobserved spatial heterogeneity such as localized cycling culture, school catchments, or perceived route safety that aggregate physical attributes cannot fully capture (Cameron and Trivedi, 2013).

To model zone-level trip generation, a Negative Binomial (Type-2 Negative Binomial (NB2)) regression is used. The conditional expectation of daily trips in zone i (μ_i) is specified using a logarithmic link function:

$$\ln(\mu_i) = \ln(N_i) + \beta_0 + \sum_{j=1}^{k} \beta_j x_{ji} \tag{3.4}$$

where $\mu_i = \mathrm{E}(Y_i \mid \mathbf{x}_i)$ represents the expected daily trip count in zone i, N_i is the population of zone i, β_0 is the intercept, x_{ji} denotes the j-th standardized explanatory attribute ($j = 1, \ldots, k$), and β_j is its corresponding parameter coefficient. The term $\ln(N_i)$ enters as an exposure offset with its coefficient constrained to 1. Rearranging Equation (3.4) yields $\ln(\mu_i/N_i) = \beta_0 + \sum_{j=1}^{k} \beta_j x_{ji}$, showing that the model predicts per-capita trip generation rates and controls for scale differences across zones. Because attributes x_{ji} are standardized, the exponentiated coefficient $\exp(\beta_j)$ represents the multiplicative factor change in expected trip rate per one-standard-deviation increase in attribute j.

To account for overdispersion, the NB2 specification expresses the conditional variance of observed daily trips Y_i as a quadratic function of the mean:

$$\mathrm{Var}(Y_i) = \mu_i + \alpha\mu_i^2 \tag{3.5}$$

where $\alpha \geq 0$ is the dispersion parameter. When $\alpha = 0$, Equation (3.5) reduces to the Poisson variance ($\mathrm{Var}(Y_i) = \mu_i$). A statistically significant estimate of $\alpha > 0$ confirms overdispersion; ignoring it would underestimate standard errors and lead to spurious statistical significance (Cameron and Trivedi, 2013).

Model parameters $\boldsymbol{\beta}$ and dispersion α are estimated jointly via maximum likelihood. Explanatory variables are selected through forward stepwise optimization based on the Akaike Information Criterion (AIC) (Akaike, 1974), stopping when no remaining candidate yields an AIC reduction greater than 2. Model validity is evaluated through three diagnostics: Variance Inflation Factor (VIF) to check for multicollinearity, Moran's I over a queen-contiguity spatial weights matrix (Moran, 1950) to test for spatial autocorrelation in model residuals, and k-fold cross-validation reporting Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) for out-of-sample predictive accuracy.

3.4.5 Population Synthesis

Trip generation returns a calibrated demand estimate $\hat{\mu}_i^{\text{cal}}$ per zone. Turning that into a workable population means producing individual agents, each carrying a complete daily

activity chain, in numbers that reproduce those zonal targets. The only source of realistic activity chains is the survey itself, so synthesis is fundamentally a resampling problem: which observed diaries are copied, how many times, and into which zone. The unit of resampling is the respondent's entire trip itinerary. Resampling trips independently would destroy the activity chains that trip distribution and assignment depend on. Two design decisions govern the outcome, and both were reconsidered for the reported model: which donors a zone may draw from, and what quantity the draw is anchored to.

Characteristic-Based Donor Matching

The obvious approach is for each zone to replicate its own respondents, falling back on a geographically nearest neighbor where a zone has none. This is straightforward but has a severe practical consequence at this sample size: each zone rests on a handful of distinct diaries, and individual respondents are replicated hundreds of times. The synthetic population then inherits the traits of a few real people as if they were regional structure. The reported model instead lets each zone draw from the whole regional pool of surveyed cyclists, giving every candidate a weight that reflects how well it suits the zone being filled. Let z denote the target zone (the one being populated) and d a candidate donor (one surveyed respondent, with a complete daily diary). The donor's own home zone is written z_d. The selection weight is the product of three factors, each capturing one notion of suitability.

$$w(d \mid z) \propto a(d,z) \cdot s(z,z_d) \cdot \tau(d,z), \tag{3.6}$$

which are defined in turn below. The weight is used only in relative terms, so the constant of proportionality does not matter: donors are drawn with probability equal to their weight divided by the sum of the weights of all candidates.

A donor who actually lives in the target zone is given a fixed multiplicative bonus, and every other donor a factor of one:

$$a(d,z) = \begin{cases} \lambda & \text{if } z_d = z \quad \text{(donor lives in the target zone)}, \\ 1 & \text{otherwise.} \end{cases} \tag{3.7}$$

The bonus is set to $\lambda = 5$ in the reported configuration, which keeps genuine local diaries prominent without making them the only ones a zone can use.

Donors from zones that resemble the target zone are preferred. Each zone is described by six standardized attributes, collected in a vector $\mathbf{x}_z$: population density, median income, cycling lane density, average household size, car ownership rate and the share of residents aged 18–65. Similarity decreases smoothly with the Euclidean distance between the two zones' attribute vectors,

$$s(z,z_d) = \exp\left(-\frac{\|\mathbf{x}_z - \mathbf{x}_{z_d}\|^2}{2h^2}\right), \tag{3.8}$$

where $\|\mathbf{x}_z - \mathbf{x}_{z_d}\|$ is that distance and h is a bandwidth controlling how quickly similarity falls off. The factor equals 1 for an identical zone and approaches 0 as the two zones diverge. This replaces geographic proximity with attribute proximity, so a peripheral zone is matched to demographically similar zones rather than to whichever zone happens to border it.

The first two factors act at the level of whole zones and cannot tell which individual donors suit the target. The third factor does so by asking, for four personal attributes, whether the donor's value is more or less common in the target zone than in the region as

a whole. Writing the four attributes as a set $\mathscr{A}$ (urbanization level, household composition, household income class, and age class), the factor is the product over them of a simple ratio,

$$\tau(d,z) = \prod_{a \in \mathscr{A}} \frac{P(a_d \mid z)}{P(a_d)}, \tag{3.9}$$

where, for attribute a, $P(a_d \mid z)$ is how common the donor's value is among residents of the target zone and $P(a_d)$ is how common it is across the whole region. A donor is up-weighted when their attributes are typical of the target zone (ratio above one) and down-weighted when they are unusual there (ratio below one). The denominator $P(a_d)$ is the reason the respondent base of Section 3.4.2 is kept in full, including non-cyclists: it is the attribute distribution of the general regional population. Estimating it from cyclists alone would conflate "typical of this zone" with "typical of people who cycle", and the ratio would no longer measure what it is meant to. Together these three factors broaden each zone from a handful of donors to roughly 60 distinct ones and cut the maximum number of times any single diary is reused from several hundred to 126.

Demand Anchoring

The second decision is how many agents a zone receives. The demand target from trip generation is a number of trips per zone, $\hat{\mu}_z^{\text{cal}}$ (Equation (3.10)), so the number of agents to draw depends on how many trips each drawn agent contributes.

The intuitive choice is to convert the trip target into a person target by dividing by an assumed average of trips per cyclist, and then draw that many agents. This is correct only if the agents actually drawn make trips at that assumed average. Under uniform copying they do, but the weighted draw of Equation (3.6) changes which diaries are selected and therefore changes their average trip count, so a person target would miss the trip target. Measured against the trip control total, anchoring on persons inflated the simulated trip total by 1.9%. The reported model anchors directly on trips. It divides the zone's trip target by the average number of trips a donor drawn under the actual weights carries, denoted $\bar{t}_z$,

$$n_z = \frac{\hat{\mu}_z^{\text{cal}}}{\bar{t}_z} \tag{3.10}$$

where $\bar{t}_z$ is computed from the same donor weights used for the draw itself, so target and realization cannot drift apart. The target n_z is generally fractional, and simply rounding it would bias the population up or down in the many zones whose target is close to a half. It is therefore realized by stochastic rounding: the whole-number part is kept, and one further agent is added with probability equal to the fractional part. For example, a target of 12.3 becomes 12 agents with probability 0.7 and 13 with probability 0.3, so that the expected number of agents equals the target exactly.

Itinerary Adjustment

A drawn agent's home zone is reassigned to the target zone for its entire itinerary, and the origin of its first chronological trip moves accordingly; the remaining stops of the day are kept where the donor respondent made them. To prevent copies of the same diary from departing simultaneously and creating artificial demand spikes, each copy receives a stochastic shift of its whole day, $\Delta t \sim \mathcal{N}(0, \sigma^2)$ with $\sigma = 30$ minutes in the reported configuration, and each activity duration is jittered uniformly by $\pm 10\%$. Departure and arrival times are then recomputed sequentially along the chain so that chronological integrity is preserved. Recomputed timelines are clamped to $[00{:}00, 23{:}59]$ rather than wrapped past midnight: wrapping produces itineraries in which an activity appears to precede its own predecessor, which the simulation in MATSim cannot interpret.

Verification Against the Survey

Because the level of the synthetic population is anchored to one weighted survey quantity, it can be checked against the other. The trip total follows from Equation (3.2) and is the anchoring target; the cyclist total follows independently from Equation (3.1) and is not targeted by anything. Agreement on both, together with a plausible implied trip rate per cyclist, is evidence that the weighting has been applied on the right basis rather than merely self-consistently. The comparison is reported with the results.

3.4.6 Trip Distribution

In the fifth part of Section 3.4, a destination is assigned to every itinerary of each person. ODiN records the purpose and the reported traveled distance of every trip, but locates its endpoints only at zone level. A traffic simulation needs coordinates. An agent must ride from one point on the network to another. Trip distribution therefore has a different task here than in a conventional four-step model. It does not estimate a zonal origin-destination matrix, which the survey already supplies in aggregate; it selects, for each individual trip of each synthetic agent, a concrete destination facility that is consistent both with the trip's purpose and with the observed distribution of trip distances. The list of destinations is a facility layer, constructed in advance by the procedure described next, in which each facility carries the activity type it can host and a size term S_f proportional to its capacity to attract that activity.

Construction of the Facility Layer

The facility layer is written to a MATSim file projected in RD New (EPSG:28992) so that facility coordinates and the network share a reference system. Two complementary open sources are combined. Facilities are extracted from OSM within the MRE boundary using `osmnx` (Boeing, 2017) across the `shop`, `amenity`, `office` and `leisure` tags. Building footprints are calculated before converting polygons to centroids, preserving size data. Building-level coverage is supplied by the BAG, which carries a registered use class and an official floor area, retrieved from the PDOK web feature service.

Both sources are mapped to the ODiN trip-purpose categories that use a facility layer (work, education, shop, leisure, personal, other, and home; see Table A.5); touring and visiting family and friends are handled separately, without a facility layer (see below). To ensure every purpose has candidate facilities, all `amenity` tags are queried.

The attractiveness term S_f in Equation (3.11) is represented by facility floor area, the registered `oppervlakte` for BAG units and the polygon footprint for areal OSM features, stored in the `capacity` attribute as a relative weight rather than an occupancy limit. Point features without footprints receive a purpose-specific fallback area (Table A.5) to preserve order-of-magnitude contrasts across activity types (e.g., 600 m^2 for schools versus 80 m^2 for corner shops), preventing S_f from collapsing into pure distance decay. Finally, the two layers are concatenated, clipped to the MRE boundary, and deduplicated on activity type and rounded coordinates to remove spatial overlaps between OSM and BAG. Floor area is a proxy for attractiveness rather than a direct measure of it, standing in for the jobs or activity scale a facility supports, so it understates open-air destinations such as parks and can equate facilities of different true attractiveness that share a footprint.

Gravity-Based Destination Choice

Each trip must be given a specific destination facility. The choice follows the intuition of the classical gravity model: a facility is a more likely destination the larger it is and the closer it is to the trip's origin, with proximity mattering more strongly for some trip purposes (a quick

errand) than others (a commute). This is formalized as a singly constrained gravity model with a negative-exponential deterrence function (Wilson, 1967). Consider a trip of purpose p (work, shop, leisure, and so on) that starts at origin o. Every facility f that can host purpose p is scored by an attractiveness value that multiplies its size by a distance penalty,

$$V_p(f \mid o) = \underbrace{S_f}_{\text{size}} \cdot \underbrace{\exp\!\left(-\beta_p\, d_{of}\right)}_{\text{distance penalty}}, \tag{3.11}$$

where S_f is the facility's floor area, d_{of} is the straight-line distance from the origin to the facility, and β_p is a purpose-specific decay rate (larger β_p means distance is penalized more sharply, so trips of that purpose stay shorter). The exponential term equals 1 at the origin and falls towards zero with distance, so a facility twice as far away is markedly less attractive unless it is correspondingly larger. The destination is then drawn at random in proportion to these scores rather than set to the single highest-scoring facility,

$$P_p(f \mid o) = \frac{V_p(f \mid o)}{\sum_{g \in \mathscr{C}_p(o)} V_p(g \mid o)}, \tag{3.12}$$

where the denominator sums the scores of all candidates $\mathscr{C}_p(o)$ within reach of the origin (the candidate set is defined below) and simply normalizes the scores into probabilities that sum to one. Drawing from this distribution reproduces the observed spread of trip distances: always taking the highest-scoring facility would send every trip of a given purpose to the same few large attractors and collapse the distance distribution onto a single value.

Calibration of the Distance Decay

The decay parameter β_p is calibrated per purpose against ODiN. For an exponential deterrence function the mean trip distance equals $1/\beta_p$, giving the initial estimate

$$\beta_p^{(0)} = \frac{1}{\bar{d}_p^{\,\text{obs}}}, \tag{3.13}$$

where $\bar{d}_p^{\,\text{obs}}$ is the observed mean trip distance for purpose p. This value is refined iteratively against the assigned mean distance, with a cap of five assignment runs. Iteration stops once no β_p changes by more than 5% between runs, and a purpose whose assigned mean distance lies within 5% of its target is not adjusted further. The reported run converged after three runs.

Because ODiN's reported distance `AfstV` is a network distance while Equation (3.11) operates on straight-line distance d_{of}, every calibration target is divided by a network detour factor $\kappa > 1$,

$$d^{\,\text{target}} = \frac{d^{\,\text{ODiN}}}{\kappa}, \tag{3.14}$$

so that assigned straight-line distances sit systematically below the raw ODiN means by $1/\kappa$; this offset is the correction working as intended, not under-assignment. As a single global multiplier, κ cannot reflect that real detour ratios vary with trip length and area type, highest for short urban trips and lowest for direct regional ones, so it under-corrects at one end of the distance distribution and over-corrects at the other.

Evaluating Equation (3.11) over every facility in the region is unnecessary, since the exponential term makes distant facilities negligible. The candidate set $\mathscr{C}_p(o)$ is restricted to facilities within a radius r_p of the origin o,

$$r_p = -\frac{\ln(1-q)}{\beta_p} \;\approx\; \frac{4.6}{\beta_p} \quad \text{for } q = 0.99, \tag{3.15}$$

where q is the fraction of the decay mass to retain. Pools exceeding a cap of 3,000 facilities are randomly subsampled to $\mathscr{C}_p(o)$. See Appendix A.5 for the convergence procedure and why subsampling is random rather than nearest-neighbor.

Purpose-Specific Treatment

Three trip purposes are not well described by a facility-based gravity model and are instead handled separately. First, for home trips, the destination is the agent's own home location, and the distance is taken as the direct distance to that location. Second, for trips made to visit family and friends, private homes lack a facility layer; destinations are therefore sampled directly from the home locations of the synthetic population, weighted by the same distance decay function in Equation (3.11) but leaving out a size term, as individual dwellings share equal intrinsic spatial attractiveness. Finally, touring trips serve recreational purposes where the ride itself is the primary objective rather than a specific destination. These trips bypass the facility model entirely: a bearing is drawn uniformly at random, and a trip distance is resampled from the empirical distribution of observed touring distances. If a sampled endpoint falls outside the study area, the distance is resampled downwards and ultimately clipped to the study boundary to ensure all generated points remain within the modeled network.

Chain Propagation and Itinerary Construction

Assigning each trip's origin independently, for instance by jittering the reported origin zone, breaks the agent's activity chain: the destination of one trip and the origin of the next end up at different points, and the simulation must silently teleport the agent between them. Origins are therefore propagated along the chain: only the first origin of the day, the agent's home, is drawn independently, and every subsequent trip starts from the destination assigned to the trip before it. The itinerary builder writes out the resulting sequence of activities and legs in MATSim plan format, taking activity durations from the survey record where available.

3.4.7 Trip Assignment

The final step of the four-step model assigns the synthetic itineraries to the cycling network, using MATSim 15 (Horni et al., 2016) with the bicycle contribution module (Ziemke et al., 2017), which adds bicycle-specific routing dis-utilities and a bicycle vehicle type to the core simulation.

MATSim does not solve a static equilibrium but runs a co-evolutionary loop: in each iteration every agent executes a daily plan in the queue-based mobility simulation (QSim), the plan is scored with a utility function, and a fraction of agents then replace their plan with a new or previously stored alternative. Repeating this drives the population towards a relaxed state in which no agent can improve its score by unilaterally changing plan. Scoring is agent-level, so the output is a disaggregated record of every link entry and exit, which is the link-level demand surface Phase 2 requires.

Scoring Function

Plans are scored with the Charypar–Nagel utility specification (Charypar and Nagel, 2005):

$$U = \sum_{q} \beta_{\text{perf}}\, t_{\text{act},q} + \sum_{a} \left(\beta_{\text{travel}}\, t_{\text{travel},a} + \beta_{\text{dist}}\, L_a + \beta_{\text{infra}}\, \phi_a \right), \tag{3.16}$$

where q indexes the activities of the plan and a the traversed links, $t_{\text{act},q}$ is time spent performing an activity, $t_{\text{travel},a}$ and L_a are the travel time and length of link a, and ϕ_a is the

link's cycling dis-utility, made up of the attractiveness term of Equation (3.3) together with the comfort, infrastructure and gradient terms supplied by the bicycle contribution from the link attributes of Section 3.4.3. Utility is gained by performing activities and lost by traveling, so an agent that reaches its activities faster or over more pleasant infrastructure scores higher.

Configuration

The router uses the link free-flow speeds of Section 3.4.3, 21.1 km/h on standard ridable links and lower on the slower link types, and the bicycle routing-speed cap of 6.03 m/s lies above all of them, so the router and the QSim agree on travel times. The run lasts 100 iterations at a learning rate of 0.1; each iteration 20% of agents re-route and 5% mutate their departure time, with innovation disabled towards the end so the population settles onto the plans it has explored. Route choice uses the attractiveness coefficient of Equation (3.3), set to -0.0008; the full parameter set is in Table A.6 of Appendix A, and the random seed is fixed and reported.

Convergence is assessed on the mean executed plan score $\bar{U}^{(n)}$ at iteration n. Writing $\mathscr{I}$ for the final five iterations, the run is accepted as relaxed when its relative variation

$$\rho = \frac{\max_{n \in \mathscr{I}} \bar{U}^{(n)} - \min_{n \in \mathscr{I}} \bar{U}^{(n)}}{\left| \operatorname{mean}_{n \in \mathscr{I}} \bar{U}^{(n)} \right|} \tag{3.17}$$

falls below 2%, a spread-based criterion because an oscillating score can have a flat trend.

Output Processing

The primary output is an events file recording every link entry and exit for every agent at second resolution. It is post-processed into the demand indicators used by Phase 2:

- **Link-level daily volume:** the number of cyclist passages per directed link over the simulated day, divided by the population scale factor where a sample rather than the full population was simulated. This is the demand surface used for CSP siting.
- **Hourly volume profiles:** the same volumes binned into one-hour intervals, which identify the commuting peaks and are used to restrict the Phase 2 demand to peak commuter trips (Section 3.5.2).
- **Origin–destination flows:** an aggregate PC4-level matrix characterising the major commuting corridors.
- **Route frequency:** how often each directed path through the network is traversed, which identifies structurally important intermediate nodes independently of absolute link volume.

Simulated hours 24–29, produced by the extended QSim end time, are folded back onto the 24-hour clock before any comparison with observed data. Clamping them into hour 23 instead inflates that hour, and discarding them loses genuine late-evening demand.

3.4.8 Model Validation

Simulated link volumes are validated against NDW bicycle counters inside the MRE, using both permanent counters and temporary counters with more than ten days of counting. No parameter of the model is fitted to these counts. They serve only to check the simulated link volumes against observed cycling volumes.

Validation Metrics

No single statistic characterises the fit of a traffic model, because volume agreement and pattern agreement are distinct properties that can move independently: a model can reproduce the regional total exactly while distributing it onto the wrong links, or track the shape of the daily profile while getting the total wrong. Geoffrey E. Havers (GEH), together with Nash–Sutcliffe efficiency (NSE) and the Pearson correlation r, are used as the primary evidence in this thesis. The remaining metrics computed alongside them for cross-checking, the aggregate ratio, Spearman's ρ, the variance ratio and Scalable Quality Value (SQV) and Kling–Gupta efficiency (KGE), are defined in Appendix A.7.

The GEH statistic is the standard measure in traffic model validation (de Dios Ortúzar and Willumsen, 2011). For a simulated volume M and an observed volume C at a given station and hour,

$$\mathrm{GEH} = \sqrt{\frac{2(M-C)^2}{M+C}}. \tag{3.18}$$

Its usefulness lies in being neither an absolute nor a relative error: it tolerates larger absolute deviations at high volumes and larger relative deviations at low volumes, which is the behavior required when comparing links whose volumes differ by orders of magnitude. Following Design Manual for Roads and Bridges (DMRB) guidance (Highways Agency, 1996), $\mathrm{GEH} < 5$ is conventionally read as a good match, $5 \leq \mathrm{GEH} < 10$ warrants closer inspection but is not necessarily wrong, and $\mathrm{GEH} \geq 10$ indicates a genuine mismatch. GEH is computed per station and hour.

The NSE (Nash and Sutcliffe, 1970) compares the model's squared error against the variance of the observations,

$$\mathrm{NSE} = 1 - \frac{\sum_i (M_i - C_i)^2}{\sum_i \left(C_i - \bar{C}\right)^2}, \tag{3.19}$$

where the sum runs over every station-hour i and $\bar{C}$ is the mean observed volume, so that $\mathrm{NSE} = 1$ is a perfect reproduction and $\mathrm{NSE} = 0$ means the model is no better than predicting the observed mean everywhere. It is the strictest of the metrics used here, because it is computed on squared errors against a baseline that already knows the correct average.

NSE is reported alongside the Pearson correlation r, which measures how well the simulated and observed series co-vary linearly, computed on the hourly series and, separately, on daily station totals; the daily figures isolate whether demand is in the right place, with the temporal profile averaged out. This pairing matters because NSE is sensitive to calibration in a way that can be mistaken for skill: applying a single scalar multiplier to every simulated value moves NSE substantially while leaving the spatial and temporal pattern unchanged. A rise in NSE is therefore evidence of improved pattern reproduction only when r rises with it; otherwise the gain is calibration or variance shrinkage (Section 3.4.8).

Interpreting NSE

NSE reaches 1 for a perfect reproduction, while $\mathrm{NSE} = 0$ means the model predicts no better than assuming every station-hour equals the overall observed mean; a useful model must therefore score well above 0, and negative values indicate it performs worse than that flat baseline. Because NSE rewards matching the overall level, it can be raised simply by rescaling all simulated volumes, without demand actually being placed on the right links. It is therefore read together with the correlation r: a high NSE with a high r means the model reproduces both the volume level and its spatial pattern, whereas a high NSE with a low r means it has only matched the average.

Acceptance Threshold

The conventional DMRB criterion (Highways Agency, 1996) is that at least 85% of comparisons achieve $\mathrm{GEH} < 5$. This threshold was developed for motorized traffic, and a fixed GEH cut-off judges low-volume counts differently from high-volume ones (Friedrich et al., 2019). It is therefore reported as a reference point rather than a strict pass or fail, and the fit is graded with the SQV quality classes used in agent-based model validation (Scherr et al., 2020). The model is assessed on the full distribution of station-level GEH and SQV values and on the aggregate ratio. Observations are filtered to Monday–Friday, matching the modeled day (Section 3.4.2), and averaged across the valid weekday observation days at each station.

3.5 Phase 2: Demand-Informed Siting of Cycle Service Points

Phase 2 translates the link-level cycling demand surface of Phase 1 into ranked, justified recommendations for CSP placement across the MRE. The central methodological decision is how demand is represented. Rather than covering residents where they live, the model intercepts the realized trajectories of individual simulated cyclists as they move along the routes they were already taking; the set of p positions that captures the most quality-weighted cycling flow is then selected by a greedy algorithm whose optimality is certified against exact and interchange solvers. This section sets out the model in six parts. Section 3.5.1 states the flow-capturing formulation and its objective and explains why CSP siting is a flow-capturing rather than a covering problem. Section 3.5.2 describes the inputs and the generation of candidate positions. Section 3.5.3 defines the eight suitability factors that grade each position. Section 3.5.4 sets the two stated weight profiles. Section 3.5.5 states the solvers and the certified optimality gap. Section 3.5.6 sets out how far the result is tested against future demand growth and against the parameters the model fixes by choice.

3.5.1 A Flow-Capturing Location Model

While flow-capturing and location-covering problems of this scale are standard and computationally easy, the optimization algorithm itself is not the primary novelty of this work. Instead, the core contribution lies in the definition of the demand object.

In conventional flow-capturing models, flows consist of origin–destination movements routed strictly along static shortest paths, forcing every trip between the same two points onto an identical route. Here, demand is instead defined as the set of realized, link-by-link trajectories of individual agents simulated in Phase 1 (3.4). Consequently, trips sharing an origin-destination pair can traverse entirely different links, as MATSim route choice dynamically responds to congestion, link attractiveness, and departure times (Section 3.4.7). Capturing behaviorally realized trajectories rather than idealized desire lines generalizes the demand representation beyond what traditional shortest-path models can achieve, directly bridging the facility location problem with the underlying behavioral simulation.

The Objective Function

Let T be the set of cycling trips reconstructed from Phase 1 that fall within the model's demand window, each trip t carrying its realized path through the network, and let Y be a chosen set of candidate positions with $|Y| = p$. The objective is the quality-weighted flow

captured by Y:

$$Z(Y) = \sum_{t \in T} q_t \max_{j \in Y} \big[\kappa(d_{tj}) Q_j \big], \tag{3.20}$$

where the symbols are as follows. q_t is the expansion weight of trip t; because this run uses the full synthetic population rather than a sample, $q_t = 1$ for every trip and the objective is a direct trip count. d_{tj} is the shortest distance from candidate position j to any point on the realized path of trip t. The capture kernel κ grades how much of a trip a position intercepts as a function of that distance,

$$\kappa(d) = \max\left(0,\, 1 - \frac{d}{R_{\text{cap}}}\right), \qquad R_{\text{cap}} = 150 \text{ m}, \tag{3.21}$$

so a position on the path ($d = 0$) captures the trip in full, a position 150 m or more away captures nothing, and the capture decays linearly in between. $Q_j \in [0,1]$ is the multi-criteria quality of position j, defined in Section 3.5.4 as the weighted aggregate of the eight suitability factors; it scales the captured flow so that the same trip intercepted at a better-sited position is worth more. The routes themselves are taken as given: the model assumes that a service point does not divert cyclists from the routes they ride, so a trip counts only if its existing path already passes within reach, and a cyclist who would make a detour to use a facility is not credited.

The same objective is solved for two demand windows, and these two windows are the two models this study reports. The baseline model, the primary case, takes T to be all cycling trips reconstructed from Phase 1 across the whole day and grades positions with equal factor weights. It answers where service points best intercept everyday cycling. The commuter model, a variant, restricts T to the morning and evening commuting peaks and grades positions with a bike-commuting weighting that emphasizes the criteria a commuting market turns on (Section 3.5.4). The two models therefore differ in which trips enter T, in the candidate set (each window keeps only the links that carry flow within it, Section 3.5.2), and in the factor weights applied. The formulation and the factor definitions are identical. Because demand is the objective and never a factor, the everyday/commuting distinction is expressed by this choice of window, not by adding a demand layer to the quality term.

The quality term enters multiplicatively. A captured trip at a poor-quality position is worth less than the same trip at a good one, and when every factor is removed the quality collapses to $Q_j = 1$ and Equation (3.20) reduces cleanly to pure flow capture. This flow-only reference (setting $Q_j = 1$ for every position) is the case against which the contribution of the multi-criteria layer is measured throughout Phase 2. It is distinct from the baseline model introduced above, which still applies the eight factors, at equal weight.

Three Properties of the Maximum

Taking the maximum per-position capture in Equation (3.20), rather than a simple sum, incorporates three desirable modeling properties without requiring explicit, ad-hoc adjustments. First, anti-clustering occurs naturally. When two candidate positions lie on the same corridor, they compete for the same trips. Because a trip captured by the first location yields almost no additional gain at an adjacent location, the model naturally distributes selected positions across distinct traffic flows without needing an explicit spatial-separation penalty. Second, the objective function exhibits monotone sub-modularity, reflecting the principle of diminishing returns. Adding a position never reduces total capture, but the marginal benefit of each new position shrinks as the selected set grows. By the classical result of Nemhauser et al. (1978), this guarantees that a greedy optimization algorithm achieves at least $(1 - 1/e) \approx 63\%$ of the optimal value, with tighter bounds certified online as described in Section 3.5.5. Finally, this approach prevents double counting. Summing traffic volumes across nearby road links counts a single trip multiple times if the route

spans several links. By taking the maximum across a trip's path, each trip contributes to the objective function only once, through the single position that captures it best.

3.5.2 Model Candidates

Phase 2 relies entirely on real, open-data sources and introduces no synthetic spatial surfaces. Aside from the Phase 1 network and simulation outputs, the required data layers are summarized in Table 3.3. Building footprints, public transit stops, cycle routes, and urbanity data enter the pipeline through suitability factors (Section 3.5.3), while the network geometry and simulated agent trajectories provide both the demand patterns and candidate positions.

Layer	Source	Used for
Network + realized trajectories	MATSim Phase 1 (full population)	demand, candidates, betweenness
Facility register (BAG-keyed)	`facilities.xml`	land-use mix, business, accessibility
Building footprints	BAG pand via PDOK WFS	open-space test
Transit stops	OSM	transit-hub factor
Cycle route network	OSM route relations	main-route factor
Urbanity	CBS Buurt 2022 (`stedelijkheid`)	urban-core factor
Study area	MRE boundary shapefile	clipping

Table 3.3: Input layers used in Phase 2, beyond the Phase 1 network and simulation output.

Demand consists of reconstructed cycling trips from Phase 1, where each trip is represented by its full path through the network. To account for different travel patterns, the model evaluates two distinct demand windows (Section 3.5.1): a baseline model capturing all-day travel, and a commuter variant restricted to peak hours (07:00–09:00 and 16:00–18:00), reflecting the primary target market for commuter-oriented service points.

Two methodological choices maintain strict consistency with the validated Phase 1 simulation. First, agent trajectories are reconstructed exactly link-by-link from the MATSim event stream, matching the model's `output_legs` record without spatial approximation. Second, the commuter time windows (07:00–09:00 and 16:00–18:00) bracket the morning and evening peaks visible in the Phase 1 temporal profile (Section 4.5). Equation (3.20) calculates total objective coverage across the corresponding trip set for each window.

Candidate Generation

A facility can only intercept trips if it sits directly adjacent to an active network link. Link adjacency is therefore treated as a structural precondition rather than a scored suitability factor. Candidate positions are placed at regular 25 m intervals along flow-carrying links and filtered using non-maximum suppression within a 25 m radius to prevent overlap. This chain-based approach replaces traditional node-based sampling, avoiding bias caused by network digitization, where dense junction clusters reflect coding artifacts rather than true spatial demand.

Candidate generation follows a simple rule-based pipeline. Directed network links are first merged into undirected corridor alignments, as opposing directions belong to the same physical corridor. The model retains only alignments carrying flow within the active demand window, places candidate points every 25 m, and thins them via non-maximum

suppression. Because the baseline window includes all-day traffic, it retains more active links and produces a larger candidate set than the commuter window.

Candidates are filtered using only one hard constraint. The underlying link must carry flow during the selected demand window. Filtering candidate sites using soft suitability criteria (such as land availability) is intentionally avoided. Using secondary factors as hard filters creates implicit, infinite-weight constraints that override the explicit weighting framework established in Section 3.5.4.

The 25 m candidate spacing provides sufficient spatial resolution to sample the suitability field without aliasing, sitting well below the 150 m capture kernel radius in Equation (3.21). For visualization, candidate scores are aggregated onto regular grids at 1000 m, 500 m, 100 m, and 25 m resolution, which allows the same surface to be read at four scales. Model optimizations are performed strictly at candidate locations; grid aggregation is purely a rendering step. To maintain computational efficiency over large study areas, grid cells are materialized only where candidate points exist. A grid cell's composite suitability equals the maximum score of its contained candidates, ensuring high-quality individual positions are not diluted by surrounding low-scoring space. Conversely, individual factor maps present the spatial mean to reflect broader neighborhood characteristics.

3.5.3 The Suitability Factors

The quality term Q_j in Equation (3.20) aggregates eight factors, each grading a candidate position based on distinct spatial attributes. These factors differ in two meaningful respects: their form (how they interpret spatial data) and their scale (the geographic reach over which they operate). Every factor is computed at the candidate resolution using real, open-source data layers.

The Five Forms

The eight factors employ one of five structural forms, summarized in Table 3.4. A point-kernel sum accumulates a distance-decayed contribution from source points within a specified reach R, each point counting once unless it carries a weight that means the same thing for every point. A point-kernel entropy measures the diversity of point types within that same reach. A zonal lookup assigns the value of the enclosing polygon to the candidate position. Finally, a line distance measures the proximity to the nearest linear feature, and a link attribute extracts a property directly from the network link on which the candidate sits.

Form	Spatial Interpretation	Factors
Point-kernel sum	$\sum$ over source points within R of weight $\cdot\, \kappa(d)$, with weight $= 1$ for B and A	X, H, B, A
Point-kernel entropy	Shannon entropy of type-shares within R	M
Zonal lookup	value of the polygon the position falls in	U
Line distance	meters to the nearest line	F
Link attribute	a property of the link the position sits on	C

Table 3.4: The five forms in which the eight factors evaluate geographic space.

Scale

Table 3.5 details the reach, kernel type, and value function used to map each raw factor output to the $[0, 1]$ scale required for Q_j. Three transformation functions are applied. A

percentile transform replaces raw values with their rank across all candidates, providing robustness against the heavy-tailed distributions typical of kernel sums. A min-max transform linearly rescales values to the target interval. A threshold cost transform maps metrics where larger values are detrimental, such as physical distances or urbanity classes, into proportionally decreasing scores. The geographic reaches span two orders of magnitude, ranging from the immediate street link to a 5000 m radius.

	Factor	Form	Reach R	Kernel	Half-wt.	Value fn
X	route junction	kernel sum	150 m	linear	75 m	percentile
C	betweenness	link attribute	the link	–	–	percentile
M	land-use mix	kernel entropy	500 m	uniform	354 m	min–max
U	urban core	zonal	Buurt	–	–	threshold (cost)
B	business density	kernel sum	1000 m	uniform	707 m	percentile
A	accessibility	kernel sum	5000 m	$\exp(-d/2429)$	2697 m	percentile
H	transit hub	kernel sum	1500 m	linear	750 m	percentile
F	main cycle net	line distance	100 m	linear on d	–	threshold (cost)

Table 3.5: The eight factors: form, reach, kernel, half-weight radius and value function. Reference scales are the candidate lattice (25 m) and the capture kernel (150 m, half-weight 75 m).

Per-Factor Definitions

To evaluate the quality term Q_j, the eight suitability factors are examined across three functional themes: network topology, land use and spatial character, and multimodal transport integration. The structural network properties comprise route junction proximity (Factor X), betweenness centrality (Factor C), and main cycle network proximity (Factor F). Land-use characteristics and local spatial demand are captured through land-use mix (Factor M), urban core classification (Factor U), business density (Factor B), and spatial accessibility (Factor A). Finally, multimodal connectivity is represented by transit hub proximity (Factor H). Each factor is detailed below to explain its mathematical formulation and spatial behavior.

The evaluation begins with route junction proximity (Factor X), which captures local network complexity. It is calculated as $X_j = \sum(\deg - 2)(1 - d/150)$ over all junctions within 150 m of a candidate position j, where $\deg$ is the undirected graph degree. By merging forward and reverse network pairs to recover this undirected degree, the $(\deg - 2)$ weighting ensures a five-way junction exerts three times the influence of a three-way intersection. Summing these values ensures that a position situated between multiple junctions naturally outscores one adjacent to a single intersection.

Operating at a much larger network scale, betweenness centrality (Factor C) applies a distance-limited node betweenness on the undirected graph using a 7500 m cutoff, corresponding to a 30-minute cycling radius. Each street link takes the mean betweenness of its two endpoints. While street centrality is an established predictor of service success (Porta et al., 2009; Sevtsuk, 2014), its inclusion carries two methodological caveats. First, it represents topological potential, which naturally correlates with the realized flow capture of the objective function. Second, it relies on a node-mean approximation due to the computational limits of true edge betweenness on large networks (Section 6.4).

Completing the network topology metrics, main cycle network proximity (Factor F) measures the distance to the nearest designated cycle route relation (OSM `network = rcn/lcn/ncn`). This is processed as a cost that equals 1 directly on a designated route and decays linearly to 0 at 100 m. By deliberately utilizing route relations rather than generic

cycleway tags, this factor successfully isolates true regional through-routes from ordinary service-road sidepaths.

Transitioning from network topology to land-use characteristics, land-use mix (Factor M) computes the Shannon entropy of facility types within a 500 m radius. This value is normalized by $\ln k$ across eight regional activity types to yield a score $M_j \in [0, 1]$ scaling from single-use to maximally mixed (Cervero and Kockelman, 1997). To avoid artificially asserting perfect single-use, candidate locations lacking any facilities within this radius are recorded as missing rather than zero-filled, as a zero would imply complete homogeneity where no data actually exists.

At the coarsest spatial scale, the urban core metric (Factor U) incorporates the official CBS urbanity class of the enclosing neighborhood. Processed as a cost factor where dense urban environments receive the highest scores, this metric is deliberately broad. All candidates within a single neighborhood share this baseline value, relying on the primary flow capture metric to distinguish specific sites within that wider geographic zone.

Targeting the commercial viability of a location, business density (Factor B) counts the 13,363 BAG-keyed work facilities within a 1000 m radius, each counted once. This density grounds the facility revenue model, which frequently relies on employer subscriptions in practice. The count is deliberately unweighted: the register's capacity attribute is the only size measure it carries and that attribute is not usable, as explained below, so B ranks positions by how many employers are nearby and not by how large they are.

Broader spatial opportunity is captured by accessibility (Factor A), evaluated as $A_j = \sum \exp(-d/2429)$ over the 59,875 BAG-keyed non-home facilities within 5000 m, each again counted once. Its exponential decay scale is strictly derived from the Phase 1 simulation configuration. By applying the modeled marginal utility of travel (-7.0 util/hr) and a standard cycling speed (17 km/h), the logit destination choice yields a decay parameter of approximately 2429 m. Consequently, accessibility decays precisely as simulated agents trade travel time against opportunity, and a facility at the candidate counts for more than one at the edge of the radius.

Both factors count facilities rather than weighting them by capacity, and the reason is a data-quality finding worth stating. The facility register carries a numeric `capacity` attribute that looks usable but does not mean one thing across the file: the OSM-derived leisure entries carry their area in square metres, the largest of them 29.2 million, and 189 entries carry 99999 or 999999 sentinel values. The non-BAG entries are 20% of the opportunities but 86% of the summed capacity, and the ten largest alone hold 37%. A capacity-weighted accessibility therefore measured proximity to large green areas rather than to opportunities: its Spearman correlation with CBS urbanity was -0.03, and only 1.6% of its top-decile positions fell in Eindhoven against 27.4% in Geldrop-Mierlo. Counting each BAG-keyed facility once raises that correlation to -0.70 and puts 88.2% of the top decile in Eindhoven. Business density had the same defect locally, where 18 work facilities at the 999999 sentinel in one cluster held 34% of the summed work capacity. An attribute that is present, numeric and plausible-looking silently encoded two different quantities, and only a spatial sanity check on where the opportunities were exposed it. The cost of the fix is stated in Section 6.4: neither factor now carries employer or facility size.

Finally, multimodal integration is evaluated through transit hub proximity (Factor H), calculated as $H_j = \sum \text{hub_weight} \cdot (1 - d/1500)$ for transit stops within a 1500 m radius. The assigned hub weights rely on OSM classifications, applying a weight of 5 to train stations, 3 to bus stations, and 1 to standard bus stops. Formulating this metric as a spatial sum ensures that major transit interchanges flanked by secondary stops outscore isolated locations, aligning with established cycle service-point siting criteria (Loendersloot, 2022). In practice, the rarity of formal bus stations in the study area reduces this factor primarily to a contrast between heavy rail stations and local bus stops.

Structural Findings on Spatial Scale

The specified reaches in Table 3.5 illustrate a functional division of labor between the suitability factors and the objective function. Factor X is the only metric with a half-weight radius (75 m) matching the flow capture kernel. All other factors operate over footprints exponentially larger. Consequently, across the 150 m neighborhood where flow capture determines exact placement, factors C, M, U, B, A, H, and F remain effectively constant. This indicates a hierarchical spatial selection process: the broad suitability factors identify the optimal neighborhood, while flow capture and route junction proximity pinpoint the specific site within it. This fine-coarse division is intentional, establishing that individual factor maps represent regional suitability preferences rather than highly localized siting precision.

A related methodological detail involves the kernel shapes. Two factors, B and M, employ a uniform kernel with a hard cutoff at their reach boundary, while X, A, and H decay smoothly. This hard cutoff reintroduces the binary coverage boundary that the flow-capture model initially sought to avoid. Harmonizing all kernels to employ smooth spatial decay remains a subject for future refinement (Section 6.4).

A final note concerns the reference cycling speeds embedded in these factors, which differ by purpose because each enters only as a distance or decay scale rather than as a travel-time prediction: the network assignment cruises at 21.1 km/h (Section 3.4.3), the accessibility decay (Factor A) is derived at 17 km/h, and the betweenness cutoff (Factor C) uses a 30-minute, 7500 m radius of roughly 15 km/h. These differences are deliberate scale choices and do not feed back into the Phase 1 assignment.

3.5.4 Multi-Criteria Weighting

The quality Q_j of a candidate is the weighted aggregate of the eight factor scores of Section 3.5.3,

$$Q_j = \sum_c w_c \, v_c(j), \qquad \sum_c w_c = 1, \quad w_c \geq 0, \tag{3.22}$$

where $v_c(j) \in [0,1]$ is the value-function score of factor c at position j and w_c is its weight, the weights being non-negative and summing to one. This additive aggregation is the standard form of Multi-Criteria Analysis (MCA), and its weights are the most contestable element of any Geographic Information System (GIS)-based application of it. Rather than elicit a single weight vector and present it as objectively correct, this study fixes the weights as two stated weight profiles, one for each model of Section 3.5.1, and reports the outcome under each. Each profile is a legible statement of a planning priority rather than an elicited optimum. This keeps the weighting inspectable and reproducible, and it exposes how sensitive the selected locations are to the priority chosen instead of hiding that sensitivity inside a single number. The two profiles are given in Table 3.6.

The equal profile gives every factor the same weight, $w_c = 1/8 = 0.125$, and is used for the baseline model. It expresses no prior view about which criterion matters most, which suits a demand window of all-day cycling: with no commuting emphasis to encode, the eight criteria enter on equal terms and it is flow capture that then decides where, within the preferred neighborhoods, the sites fall. The bike-commuting profile is used for the commuter model and doubles three factors that a commuting service market turns on, namely the urban core (U), business density (B) and main cycle network proximity (F). Against a base weight of one for the remaining factors, doubling these three and renormalizing gives them $2/11 \approx 0.182$ each and the other five $1/11 \approx 0.091$ each, so the profile is a simple, legible statement of emphasis rather than a finely tuned vector. The reasoning is that a commuting-oriented service point should sit in denser urban settings (U), near the employer base that underwrites the subscription business model (B), and on the designated through-routes that carry commuting cyclists (F).

A structural caveat applies to any additive multi-criteria scheme. A weight expresses the importance assigned to a criterion, not the statistical independence of that criterion from the others, so weighting several correlated factors separately still counts a shared underlying dimension more than once. The overlap between the eight factors is therefore measured before their weighted outcome is interpreted (Section 5.3.1).

	Factor	Equal	Bike-commuting
X	route junction	0.125	0.091
C	betweenness	0.125	0.091
M	land-use mix	0.125	0.091
U	urban core	0.125	**0.182**
B	business density	0.125	**0.182**
A	accessibility	0.125	0.091
H	transit hub	0.125	0.091
F	main cycle network	0.125	**0.182**
	sum	1.000	1.000

Table 3.6: The two stated weight profiles. The baseline model uses the equal profile and the commuter model the bike-commuting profile, which doubles the urban-core (U), business-density (B) and main-cycle-network (F) factors before renormalizing. Weights sum to one.

Discussion of the Weights with Practitioners

The profiles were not left as a private judgement. They were put to a mobility expert at the Province of Noord-Brabant, together with the reasoning behind the three doubled factors and the list of eight criteria itself. Three points came out of that discussion and are recorded here because they bound what the weighting can claim.

The expert judged the criteria and the emphasis of the bike-commuting profile sound for a commuter-oriented facility, which is an outside check on the face validity of the set. The expert also proposed alternative emphases, in particular a heavier weight on transit interchange, which were not run and remain open. And the discussion confirmed that practitioners state importance rather than independence: asked which criteria matter, none of the parties to the conversation raised the question of whether two criteria measure the same underlying property, which is the structural weakness any additive scheme carries (Section 5.3.1).

A single conversation is not an elicitation. It gives the profiles a practitioner reading without the pairwise comparisons, the consistency indicator or the aggregation across respondents that a structured method such as the Analytic Hierarchy Process or the Best–Worst Method would supply (Section 2.6.4). The weights therefore remain declared policy scenarios rather than validated expert consensus, and neither profile is claimed as objectively correct. What makes this acceptable for the present purpose is an empirical result rather than an argument: the sensitivity analysis of Section 3.5.6 measures how much the weighting can move the answer at all, and reports it alongside the answer. The limitation is restated in Section 6.4.

3.5.5 Solving the Location Problem

The facility selection process maximises quality-weighted flow capture subject to a fixed candidate budget:

$$\max_{Y} \; Z(Y) \quad \text{s.t.} \quad |Y| = p, \tag{3.23}$$

where Y is the selected subset of candidate positions and p is the facility budget. To provide a flexible implementation ladder rather than a single fixed outcome, Equation (3.23) is solved across budget steps $p \in \{3, 5, 7, 10\}$; these steps are chosen for illustration and can be changed to whatever number of locations is required.

The problem is solved with a deterministic greedy algorithm. Starting from an empty set, at each of the p iterations it adds the candidate with the largest marginal gain $\Delta(j \mid S) = Z(S \cup \{j\}) - Z(S)$, where S is the set chosen so far. Because the objective is monotone submodular (Section 3.5.1), this greedy rule carries a worst-case guarantee: its solution captures at least $(1-1/e) \approx 63\%$ of the theoretical optimum (Nemhauser et al., 1978). Greedy is the standard solver for a submodular coverage objective of this kind, and the order in which it adds positions also gives each chosen site a natural rank, best first, which is used when the two models are compared below.

Certifying Solution Optimality

The $(1-1/e)$ figure is only a worst-case floor, so the quality of the actual solution is certified per instance with an online submodular upper bound (Minoux, 1978; Leskovec et al., 2007) rather than by solving the problem exactly. At any greedy iteration with selected set S, the true optimum Z^* is bounded above by the current objective plus the sum of the p largest marginal gains still available:

$$Z^* \leq Z(S) + \sum_{k=1}^{p} \Delta_{(k)}(S), \tag{3.24}$$

where $\Delta_{(1)}(S) \geq \cdots \geq \Delta_{(p)}(S)$ are the ordered marginal gains across the remaining candidates. Computing this bound costs only one extra pass over the capture matrix and yields a computable, instance-specific optimality certificate that is tighter than the conservative worst-case guarantee; the concrete gap values are reported with the results.

Model Comparison via Site Agreement

To evaluate how the demand window affects facility placement, the baseline (all-day) and commuter (peak-period) models are solved independently across the budget ladder, and their selected candidates are ranked by greedy selection order. The spatial overlap between the two is then measured by site agreement, the number of top-p positions that coincide on the same or an adjacent candidate. Site agreement shows whether changing the demand window materially relocates the recommended sites or merely reconfirms the same high-flow corridors that flow capture would select in either case.

Reading the Selected Sites

Four measures describe the structure of a solution, and the results report them without restating the method. First, a site is invariant when it is selected at every budget $p = 3, 5, 7, 10$, which marks the positions that any solution of any size includes. Second, the own flow of a selected site j is its leave-one-out contribution $Z(Y) - Z(Y \setminus \{j\})$, the captured flow lost if that site alone were dropped. Because nearby sites can capture the same trips, own flows need not decrease with selection rank and sum to less than the total coverage. Third, the best rival of a selected site is the unselected candidate within 150 m that would add the most flow in its place, evaluated against the same remaining set of sites, and the number of alternatives is the count of unselected candidates within that radius. A rival close by with nearly the same gain shows that the exact position is one of several near-equal points on the same corridor.

Fourth, to show where each factor on its own would place sites, a single-factor siting ranks all candidates on that factor alone, ignoring flow, and selects p of them

by non-maximum suppression: it takes the highest-scoring candidate, excludes every candidate within a separation radius of it, and repeats. Without this step all p sites would cluster on the single highest-scoring spot. The radius is set to 0.45 times the spacing of an even grid of p points over the bounding box of the study area, which gives 10 km at $p = 5$, so the single-factor sites are spread across the region by construction. The flow those sites would capture is then compared with the flow captured by siting on demand alone.

3.5.6 Robustness of the Siting Result

Several quantities in the model are set by choice rather than estimated: the factor weights, the demand window, the level and composition of future demand, the capture radius, the reach of the factor buffers, the candidate spacing, and the decision not to impose a minimum separation between selected sites. Four analyses test how far the siting result depends on them.

Each analysis reports two quantities, the flow captured at the reporting budget $p = 5$ and how many of the five selected positions are shared with the unperturbed run. A perturbation can move coverage without moving sites, and the planning question turns on the sites. The capture radius is 150 m unless it is itself the parameter being varied. The scenario analysis starts from the two reported models, and the weight, structural and separation analyses all start from the baseline model under its own equal weights.

Current and Future Demand Scenarios

The scenario analysis crosses three current scenarios with four future growth cases. A current scenario is a pair of a factor-weight set and a demand model (Table 3.7). S1 is the baseline model, all-day demand under equal weights. S2 is the commuter model, peak-hour demand under the bike-commuting weights of Table 3.6, run on its own set of $235{,}844$ candidates rather than the $269{,}379$ of the all-day demand. S3 combines all-day demand with objective weights from the CRITIC method (Diakoulaki et al., 1995), which derives the weights from the factor data instead of from judgement,

$$C_c = \sigma_c \sum_k (1 - r_{ck}), \qquad w_c = \frac{C_c}{\sum_k C_k}, \tag{3.25}$$

where σ_c is the standard deviation of the scores of factor c across the all-day candidates and r_{ck} is its correlation with factor k. A factor gains weight when it separates candidates well and when it carries information the other factors do not. The weight set changes only the quality Q_j of Equation (3.22) and leaves demand untouched, so S1 and S3 differ in weights alone, while S1 and S2 differ in both weights and demand. The S1 and S2 solves are checked against the reported models and must reproduce their coverage and their five sites at $p = 5$.

The factor scores are rank-scaled, so their standard deviations are nearly equal and the CRITIC weights are set mainly by redundancy. Main cycle network (F) is the only factor that correlates weakly with every other factor, and it receives twice an equal share, while the seven correlated urbanness factors are held between 0.09 and 0.13. CRITIC is therefore the one weighting used here that responds to the overlap an expert elicitation cannot see, because an expert states importance and not independence. The entropy weight method was also considered and rejected, because on rank-scaled scores it responds to ties and zero mass rather than to information and places two thirds of the weight on urban core and main cycle network alone.

Future growth is applied by reweighting the trips Phase 1 already produced rather than by re-simulating. The coverage reported by the model is a share, captured flow over total flow, so a uniform multiplier on every trip weight cancels in the ratio and leaves both the coverage and the greedy choice of sites unchanged. Only a reweighting that differs between

	Factor	S1 (all-day)	S2 (peak)	S3 (all-day)
X	route junction	0.125	0.091	0.115
C	betweenness	0.125	0.091	0.127
M	land-use mix	0.125	0.091	0.108
U	urban core	0.125	0.182	0.105
B	business density	0.125	0.182	0.100
A	accessibility	0.125	0.091	0.101
H	transit hub	0.125	0.091	0.092
F	main cycle network	0.125	0.182	0.252

Table 3.7: Factor weights of the three current scenarios, normalized to sum to one. S1 and S2 run the baseline and commuter models, and S3 runs all-day demand under the CRITIC weights of Equation (3.25).

trips can move the result. Growth is therefore targeted at one trip-length band at a time,

$$q_t \;\rightarrow\; q_t\, m(t), \qquad m(t) \;=\; \begin{cases} 1+L & \text{if trip } t \text{ lies in the targeted band,} \\ 1 & \text{otherwise,} \end{cases} \tag{3.26}$$

where the band is either short trips, with a routed path shorter than 7.5 km, or long trips, of 7.5 km and more, and L is the growth level of that band. Each current scenario is run at four futures, short $+25\%$, short $+50\%$, long $+25\%$ and long $+50\%$, and the same greedy solver is applied to each.

The 7.5 km boundary corresponds to 30 minutes of cycling at an average speed, which is the range used to delimit short trips in Dutch research on the choice between active modes and the car (Scheepers et al., 2013). It also equals the betweenness cutoff of factor C, so the split is consistent with the reach already assumed inside the model. The two bands represent different kinds of growth. Short trips form the bulk of present cycling, whereas the e-bike is expected to replace the conventional bicycle on shorter distances and the car on longer ones (Huang et al., 2024), so growth in the long band stands for cycling taking over longer, car-competitive trips.

The levels of 25% and 50% are exploratory magnitudes chosen to show the direction and size of the response, not forecasts. For scale, a national projection fitted to the Dutch national travel survey expects total cycled distance in the Netherlands to grow by 13% between 2022 and 2028, driven by a 43% growth in e-bike distance against a 4% decline for the conventional bicycle (Huang et al., 2024). Because a band holds only part of all trips, the growth in total demand implied by a band level is smaller than the level itself.

Weight Sensitivity

The factor weights are the most contestable input of Section 3.5.4, so their effect on the outcome is measured directly. Starting from the equal profile, one weight at a time is multiplied by 0.5 and by 2, the weight vector is renormalized to sum to one, the quality Q_j of every candidate is recomputed through Equation (3.22), and the problem is re-solved. The resulting spread of coverage across the eight factors is a tornado diagram on the outcome.

The tornado moves one weight at a time. Joint uncertainty in all eight weights is tested by drawing 100 weight vectors from a Dirichlet distribution centered on the equal profile, $\mathbf{w} \sim \mathrm{Dirichlet}(\alpha\, w_{\mathrm{eq}})$ with $w_{\mathrm{eq}} = 1/8$ and concentration $\alpha = 50$, and solving the baseline model under each draw. At this concentration a single weight ranges from about 0.03 to 0.29 across the draws. The result is the distribution of coverage and how often each position is selected.

A weight and its influence are not the same thing. A factor that scores almost every candidate alike contributes nothing to the ranking however heavily it is weighted, so the

nominal weight w_c is also compared against the realized influence $w_c\,\sigma_c$, the weight scaled by the standard deviation of that factor's scores across the candidate set.

Structural Sensitivity

Three structural parameters are swept around their chosen values: the capture radius R_{cap} of Equation (3.21) at 75, 150, 300 and 500 m, all factor buffer radii jointly at 0.5, 1 and 2 times their stated values, and the candidate spacing at 25, 50 and 100 m. Each point is a full rebuild of the affected part of the chain, so a change in candidate spacing regenerates the candidates, the factor scores, the capture matrix and the quality surface before the problem is re-solved.

The accessibility factor is excluded from the buffer sweep. Its effective reach is set by the decay scale derived from Phase 1 rather than by its cutoff radius, at which the decayed weight is already near zero, so scaling the cutoff barely moves the factor while making the underlying spatial join intractable.

Minimum Separation Between Sites

Section 3.5.1 argues that the maximum in the objective makes anti-clustering intrinsic, so that no explicit spatial-separation penalty is needed. This claim is tested in two ways. The first is diagnostic and changes nothing in the model: the realized pairwise distances between the selected sites are measured. The second imposes the mechanism the argument says is unnecessary. A hard minimum separation D is enforced inside the greedy rule, so that any candidate within D of an already selected site is excluded for the rest of the run, and D is swept from 0 to 7,500 m. The coverage lost at each D is the price of forcing the sites apart.

3.6 Ethical Considerations

This study adheres to the ethical guidelines of Eindhoven University of Technology. The ODiN microdata are anonymised at source by CBS and are accessed through the DANS archive under its conditions of use. No re-identification is attempted and no individual record is reported. The synthetic agents produced in Section 3.4.5 are statistical replicas, not survey respondents. Home locations are drawn within a PC4 zone rather than taken from any real address. OSM, CBS and NDW data are publicly available. All intermediate data are stored on university infrastructure and are reported only in aggregated form.

3.7 Conclusion

This chapter has set out the complete methodological framework for demand-informed CSP location planning. Phase 1 addresses the data scarcity that characterizes metropolitan-scale cycling research. A bicycle-only network is constructed from OSM. A NB2 regression distributes cycling demand across zones while its level is anchored to a directly weighted survey control total. Characteristic-based donor matching expands a few hundred real diaries into a full synthetic population. A gravity-based destination choice model converts zone-level trips into concrete, chained itineraries. A MATSim model assigns those itineraries to the network until the system converges. The result is a link-level cycling demand surface covering every corridor in the MRE, not only those where empirical counts exist, together with an explicit statement of how well that surface reproduces the counts that do exist.

Two properties of this design deserve emphasis, because they bound what the demand surface can be used for. The level of demand and its spatial distribution are established by

different mechanisms and are therefore reliable to different degrees. The level rests on survey weights and is accurate, while the spatial distribution rests on a weak regression over zones coarser than a typical cycling trip. And the validation is deliberately multi-metric, because the properties that matter for CSP siting, namely whether demand is concentrated on the right corridors, are not the properties that a single aggregate volume statistic measures.

Phase 2 translates the demand surface into actionable recommendations by treating siting as a flow-capturing rather than a covering problem. Candidate positions along the cycle network are graded on eight open-data criteria under two stated weight profiles, and the set that captures the most quality-weighted cycling flow is selected by a greedy algorithm with a certified optimality gap. The demand object, and not the optimizer, carries the contribution. Demand is the realized link-by-link trajectory of each agent, so identical origin-destination pairs may contribute different paths, a representation a conventional flow-capturing model cannot make.

Because every component relies exclusively on nationally available open data and open-source software, the framework transfers to other Dutch urban regions with minimal adaptation. The following chapter presents the results of applying it to the MRE.

Chapter 4

Phase 1 Results: A Validated Cycling-Demand Surface for the MRE

This chapter applies Phase 1 of the framework (Section 3.4) to the Metropolitan Region Eindhoven (MRE). It reports, in the same order as the methodology, the trip generation model and its expansion to the region, the synthetic population and its calibration, the destination-choice distribution, the network assignment and the resulting link-level demand surface, and the validation against observed counts. The link-level demand surface produced here is the input to Phase 2 (Chapter 5).

4.1 Trip Generation

Phase 1 begins with the zonal cycling-trip production model of Section 3.4. The Type-2 Negative Binomial (NB2) model was estimated on the 2022 Onderweg in Nederland (ODiN) wave, restricted to Monday–Friday travel, and used to expand a zonal control total to the full population of the MRE. Population enters the model as an offset and therefore carries the trip level directly, while the two retained covariates set only the gradient in the per-capita cycling rate around it. The absolute demand level is anchored to an independent survey control total.

4.1.1 Model Estimation

The seven candidates cover residential density, the socio-economic make-up of each zone and the local cycling-network supply. Their spread across the 113 estimation zones and their collinearity are summarized in Table 4.1.

Figure 4.1 shows the pairwise correlation between the candidate predictors. Each cell gives the Pearson correlation of two predictors, from dark for a strong negative to pale yellow for a strong positive correlation. The strongest correlations all involve population density, which correlates 0.87 with cycling-lane density, -0.80 with car ownership and 0.75 with employment density, and this is why it carries the highest Variance Inflation Factor (VIF) and is the predictor dropped. The two predictors the model keeps, land-use mix and cycling-lane density, are themselves only moderately correlated (-0.54), and every retained candidate sits well below the conventional VIF threshold of 10 (the highest retained value is 5.46). Figure 4.2 shows the survey coverage behind the estimation sample. The shading gives the number of 2022 ODiN respondents per zone, from one to seventeen, and it

Predictor	Unit	Mean	SD	Min	Max	VIF
pop_density†	inh/km²	1,989	2,131	47	9,097	10.22
cycling_lane_density*	m/km²	26,506	11,935	8,524	50,893	5.46
employment_density	jobs/km²	252	432	7	3,916	5.33
car_ownership_rate	cars/hh	1.15	0.26	0.43	1.75	3.01
land_use_mix*	index	0.44	0.10	0.23	0.73	2.52
share_18_65	fraction	0.61	0.05	0.50	0.90	2.37
median_income	EUR/yr	30,077	4,587	22,881	42,743	1.40

† dropped for collinearity, * retained by the Akaike Information Criterion (AIC) search.

Table 4.1: The seven candidate predictors: their spread across the 113 estimation zones and their variance inflation factor, computed before selection.

follows the zone's population (Spearman $\rho = 0.79$) rather than a clean pattern from core to periphery. The 26 hatched zones have no usable observation and take a borrowed expansion factor. The takeaway is that every zone rests on a handful of respondents at most.

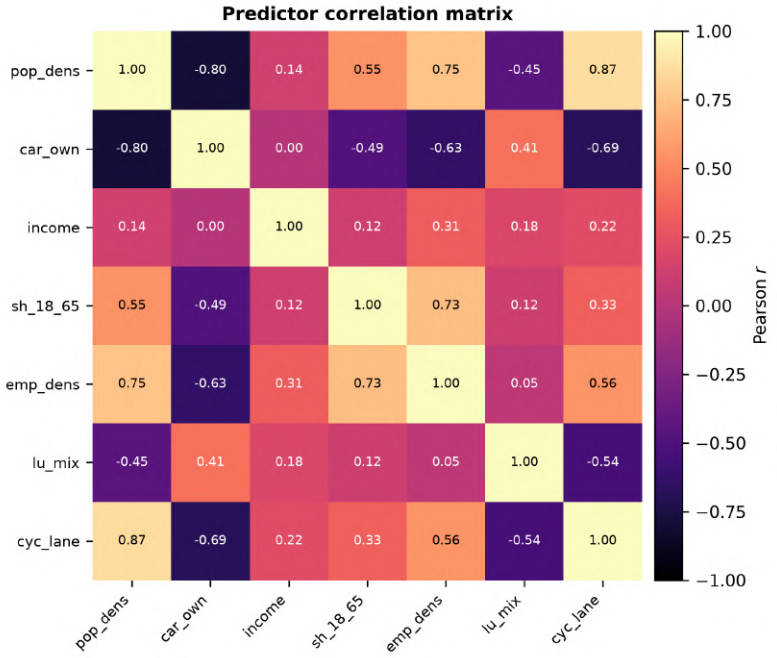

Figure 4.1: Pearson correlation between the seven candidate predictors across the 113 estimation zones.

Of the seven candidate predictors, pop_density was removed first for collinearity: its VIF of 10.2 was above the threshold of 10, while every retained candidate sat below 5.5 (Table 4.1, full diagnostics in Table A.2). Forward stepwise AIC selection over the remaining predictors kept two, land_use_mix and cycling_lane_density. No further predictor lowered the AIC. The estimated model is reported in Table 4.2.

Both retained coefficients are positive and significant at the 5% level. A one-standard-deviation rise in land-use mix raises the expected per-capita cycling rate by 21%, and the same rise in cycling-lane density by 15%. The overdispersion parameter is $\hat{\alpha} = 0.314$ and the null $\alpha = 0$ is clearly rejected ($z = 7.88$), so the negative binomial form is needed: a Poisson model would

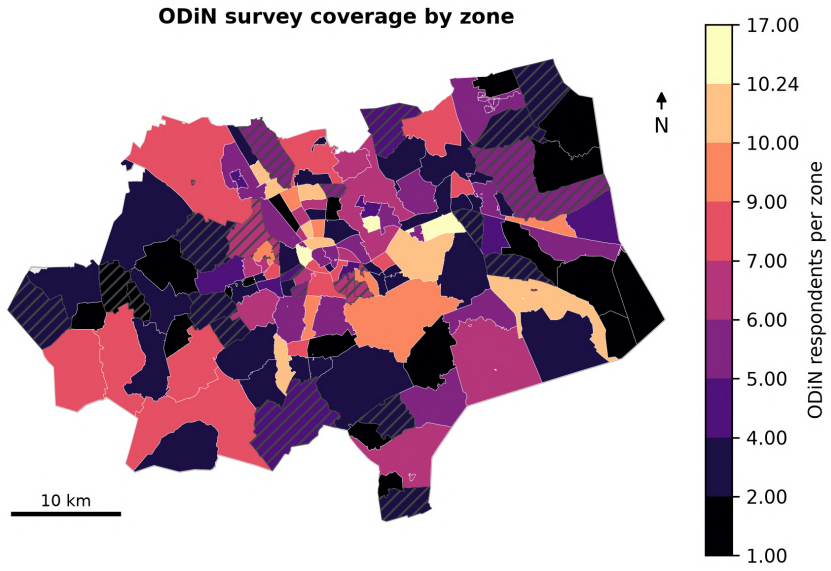

Figure 4.2: Number of 2022 ODiN respondents per 4-digit postcode zone (PC4) zone. Hatched zones carry no usable observation and take a borrowed expansion factor.

Term	β	Std. err.	z	p	$\exp(\beta)$
Constant	−0.3165	0.053	−6.00	< 0.001	–
`land_use_mix`	0.1922	0.061	3.15	0.002	1.212
`cycling_lane_density`	0.1428	0.064	2.24	0.025	1.154
Dispersion α	0.3139	0.040	7.88	< 0.001	–

Table 4.2: Estimated NB2 trip-generation model. 2022 ODiN Monday–Friday data, $n = 113$ observed PC4 zones, offset $\ln$(population). Predictors are standardized, so $\exp(\beta)$ is the multiplicative effect on the per-capita cycling rate of a one-standard-deviation increase.

understate the standard errors.

4.1.2 Fit and Predictive Validity

Because population enters as an offset, the fit is judged against two nested reference models rather than one, and each comparison answers a different question. Both are likelihood-ratio tests, in which twice the difference in log-likelihood is distributed as χ^2 with degrees of freedom equal to the number of restricted parameters. Table 4.3 reports both.

The first reference is the null model: an intercept only, with no population offset and no covariates, which predicts the same trip count for every zone regardless of how many people live in it. Its log-likelihood is -1045.6 against the full model's -1015.4. The null hypothesis that the offset and both covariates are jointly zero is rejected, $\chi^2(3) = 60.3$, $p < 0.001$, with a McFadden pseudo R^2 of 0.029. This test establishes only that the model beats predicting a constant, and most of what it rejects is the absurdity of a zone's cycling being unrelated to its population: the offset alone accounts for 24.9 of the 30.2 log-likelihood points.

The second reference is therefore the one that tests the covariates. Against a population-only model, the null hypothesis that the coefficients on land-use mix and cycling-lane density are both zero is rejected at the 1% level, $\chi^2(2) = 10.6$, $p = 0.005$, with a pseudo R^2 of 0.0052. The two covariates do carry signal beyond exposure, and that is the claim the model relies on, since population is given rather than estimated.

Neither pseudo R^2 should be read as the R^2 of a linear regression. McFadden's measure compares log-likelihoods rather than explained variance, and for count data it is not on the scale of the linear R^2 (Cameron and Trivedi, 2013). The 0.0052 is additionally an incremental value, measured over a model that already contains the dominant predictor, so it is the gain the covariates add and not the share of cycling the model explains.

How Accurate Are the Predictions?

Significance and accuracy are different properties, and this model has the first without much of the second. Three diagnostics say so, and they agree.

The coefficients are precise enough to be distinguished from zero but small in effect. A one-standard-deviation rise in land-use mix moves the expected per-capita rate by 21% and the same rise in cycling-lane density by 15%, against zonal rates that vary by considerably more than that across the region, so the covariates reorder zones only modestly around the level population sets.

Out-of-sample accuracy is weak and unstable. Across the six cross-validation folds the out-of-sample R^2 ranges from -0.80 to $+0.57$ and is negative on average, meaning that on a held-out fold the fitted model is often a worse predictor of a zone's cycling than the sample mean. With 113 zones split six ways, each fold withholds fewer than 20 zones, and a single unusual zone moves the fold statistic a long way, which is why the range is as wide as it is.

The scatter of predicted against observed zonal trips (Figure 4.3), in which each point is one estimation zone and the diagonal marks a perfect prediction, shows the same thing directly: the points follow the $1{:}1$ line in the aggregate but with wide spread. The model reproduces the general level of a zone's cycling without tracking its zone-to-zone variation.

The consequence for the rest of Phase 1 is specific rather than general. Zonal predictions are used to distribute a regional total, not to stand on their own, and the total itself is fixed by an independent control total (Section 4.1.3) rather than by the regression. An error in a single zone therefore shifts agents between zones but cannot change how much cycling the region contains. The corridor pattern that Phase 2 depends on is produced further downstream, by routing those agents over the network, and is less exposed to this scatter than the zone totals are.

The cause is the thin survey base rather than the specification. Only 113 of the 139 MRE zones carry any direct 2022 observation, and the 2022 wave contributes just $1{,}476$ weighted

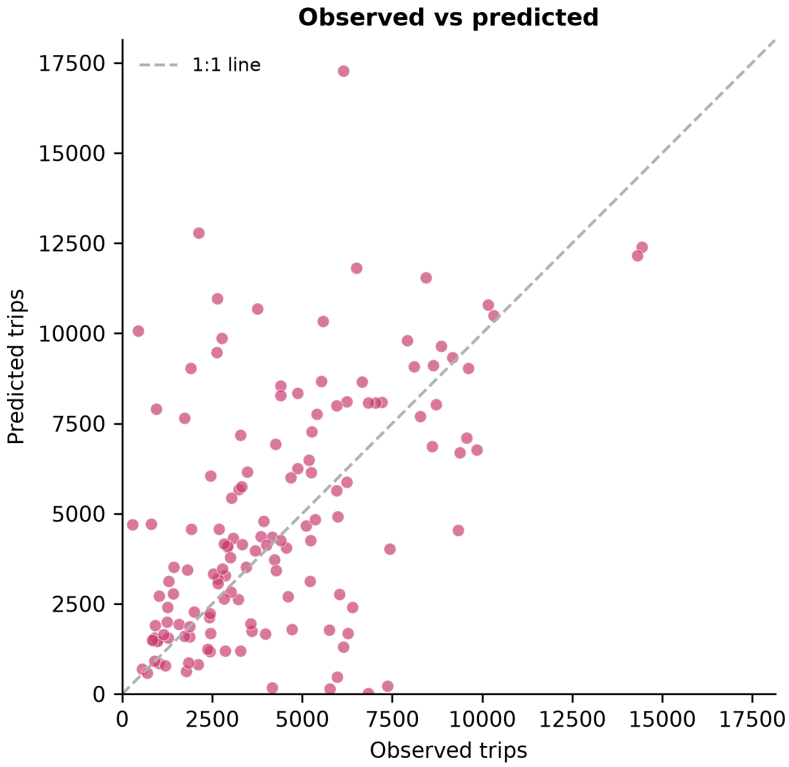

Figure 4.3: Observed against predicted zonal cycling trips (NB2 model) with the 1:1 line.

Monday–Friday bicycle trip records across them, so the per-capita rate rests on a handful of diaries per zone.

4.1.3 Expansion to the Region

The fitted rates were turned into person-level expansion factors for all 139 MRE zones. The 26 zones without a usable 2022 observation took their factor by nearest-neighbor borrowing from observed zones of similar character. The expanded demand was then rescaled to the ODiN `FactorV` control total for average-weekday bicycle travel by MRE residents, a step that moves the total from 579,082 to 680,139 trips per day (a factor of 1.175). The rescale is an increase because the regression fixes the per-capita gradient, not the absolute level: summed over the full population the fitted rates fall about 15% short of the survey's `FactorV` total, and the anchoring supplies that level rather than correcting the regression.

At this control total there are 680,139 bicycle trips on an average weekday, meaning an average Monday-to-Friday day, which is the basis every daily figure in this thesis is reported on. That is 0.86 bicycle trips per inhabitant per day, inside the Dutch metropolitan range of 0.7 to 0.9. These totals feed the synthetic population (Section 4.2), which realizes 254,010 cycling agents, a 32.1% participation rate that matches the ODiN survey (Section 4.2.2). The per-person expansion factor is mapped in Figure 4.4, where each zone is colored by the number of residents one survey respondent stands for, in quantile classes, and brighter means more. It is spatially patchy rather than a smooth gradient: most zones sit in a middle band, and a scattered handful carry much larger factors, the zones where a few survey respondents stand in for many residents. The three highest factors, between 1,783 and 5,037, belong to two residential zones in Woensel-Noord in Eindhoven (PC4 5632 and 5626) and one in Valkenswaard (5554), each represented by only one or two survey respondents.

Statistic	Value
Zones used in estimation	113
Zones predicted	139
Model degrees of freedom	2
Log-likelihood	
Null model (intercept only, no population)	−1045.6
Population offset and intercept	−1020.7
Full model (offset and two covariates)	−1015.4
Full model against the null model	
Likelihood-ratio statistic $\chi^2(3)$	60.3
p	< 0.001
McFadden pseudo R^2	0.029
Full model against the population-only model	
Likelihood-ratio statistic $\chi^2(2)$	10.6
p	0.005
McFadden pseudo R^2	0.0052
Six-fold cross-validation	
Out-of-sample R^2, range	−0.80 to +0.57

Table 4.3: Fit statistics of the estimated NB2 model and the two nested models it is tested against. The null model has neither the population offset nor the covariates and predicts one constant trip count for every zone.

They mark where the survey is thinnest rather than where cycling is highest.

Figure 4.5 shows the distribution of the per-person expansion factor: most zones sit in a narrow band while a small number carry much larger factors, the right-skew that comes from expanding a thin survey sample up to the full population. Figure 4.6 maps the resulting cycling trips per inhabitant, again in quantile classes. Most zones lie close to the regional value: 65% fall between 0.79 and 1.10 trips per inhabitant. The highest values, above 2, occur in zones with very few residents, such as the Flight Forum business area near Eindhoven Airport (PC4 5657), the Ekkersrijt business park in Son (5692) and the Oirschotse Heide (5689, seven residents). Three of the four highest zones take a borrowed expansion factor, so there the ratio divides the trips of a few residents by a very small population and says little about how much cycling takes place in these areas.

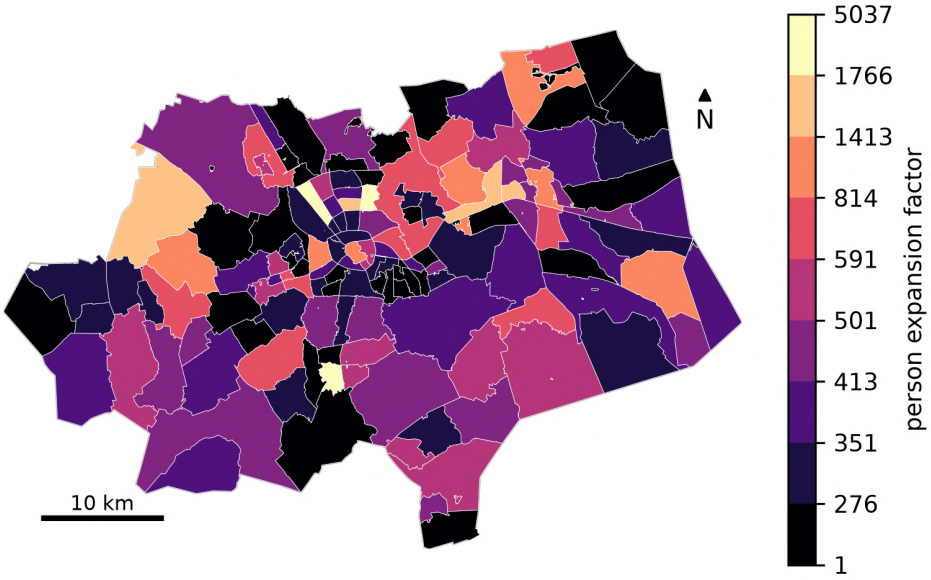

Figure 4.4: Person-level expansion factor by PC4 zone across the MRE.

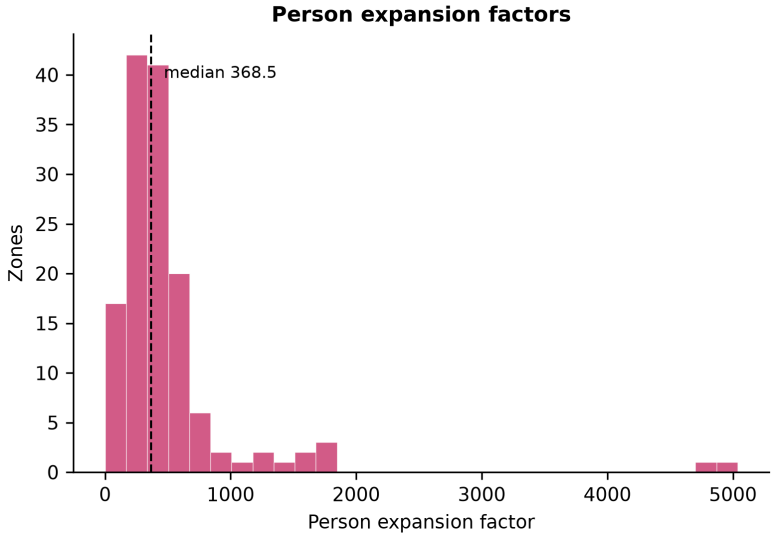

Figure 4.5: Distribution of the person-level expansion factor across the 139 MRE zones.

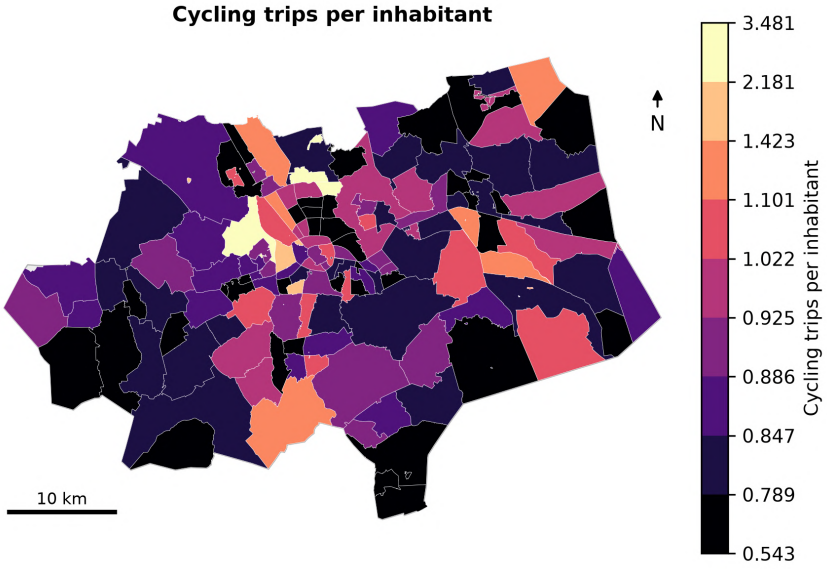

Figure 4.6: Simulated cycling trips per inhabitant by PC4 zone across the MRE.

4.2 Synthetic Population and Demand Calibration

Applying the donor matching and trip anchoring of Section 3.4 to the calibrated zonal demand produces a synthetic population of 254,010 cycling agents carrying 679,448 bicycle trips on an average weekday, at full population scale.

4.2.1 Donor Diversity

Because only 547 real Monday–Friday cyclist diaries underlie the 2022 dataset, there is a risk that a synthetic population of a quarter of a million agents simply copies a handful of respondents. The characteristic-based donor matching guards against this: each zone draws from the whole regional cyclist pool, weighted by how similar each donor is in its attributes, so a zone uses 231 distinct donors on average (median 230, range 5 to 413, Figure 4.7). The zones spread broadly around this average, marked by the dashed line, from roughly 100 to 360 donors, with only a few drawing fewer than 50. This does not remove the underlying thinness of the sample, which is recorded as a limitation, but it prevents the spatial clustering that plain copying would produce.

Figure 4.8 shows how much of each zone's population is drawn from its own donors rather than from the wider regional pool. For most zones this own-zone share is small, on average 23%, so the bulk of each zone's agents come from attribute-similar donors elsewhere in the region, which is what keeps the sample from cloning a few local diaries. Figure 4.9 maps the resulting number of synthetic agents per zone, which follows the residential population and so concentrates in the urban cores.

4.2.2 Calibration Against the Survey Weights

The demand level is anchored to the ODiN `FactorV` trip control total, so agreement on the trip count is true by construction and is not evidence of fit. The value of the calibration lies in the quantities that are not targeted. Table 4.4 compares the synthetic population against the two independent ODiN weights, `FactorV` (trips) and `FactorP` (persons), in which only the trip total is a calibration target and the remaining rows are independent checks.

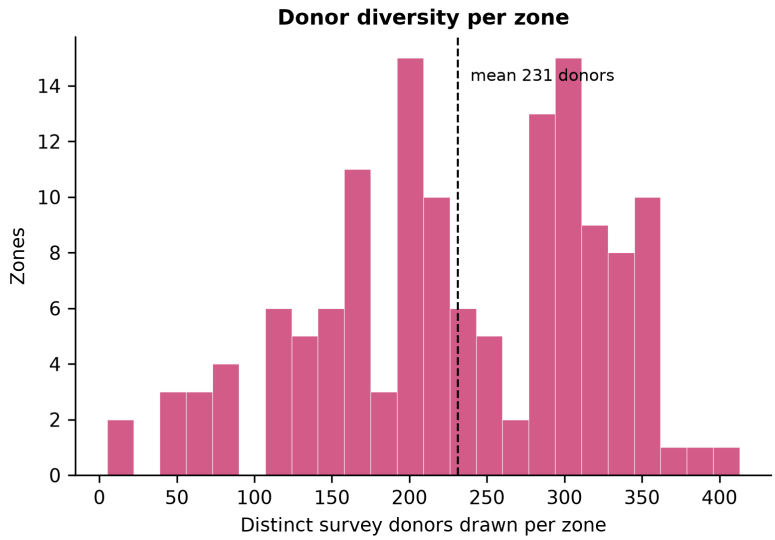

Figure 4.7: Number of distinct survey donors drawn per PC4 zone across the 139 MRE zones. The dashed line marks the mean.

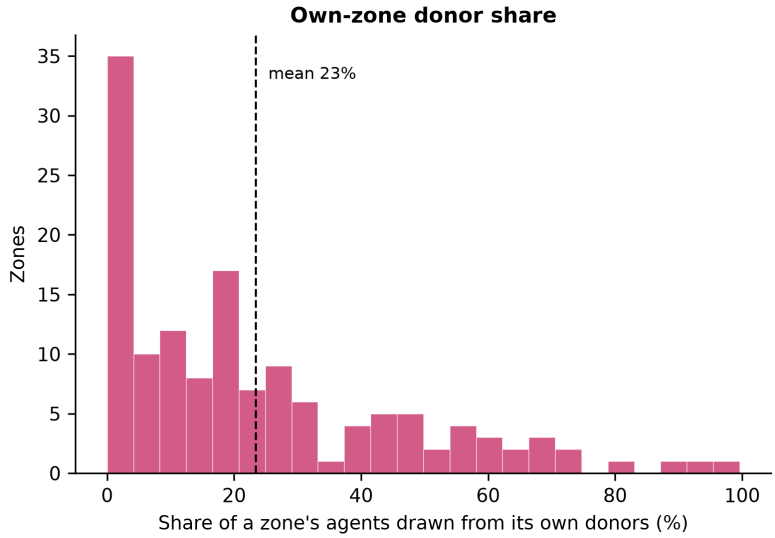

Figure 4.8: Share of a zone's agents drawn from donors living in that same zone, across the 139 MRE zones.

Quantity	ODiN	Simulation	Ratio
Trips per weekday	680,139	679,448	0.999
Cyclists per weekday	252,265	254,010	1.007
Trips per cyclist	2.696	2.675	0.992
Cycling participation	31.9%	32.1%	–

Table 4.4: The synthetic population compared with the ODiN survey weights on an average weekday.

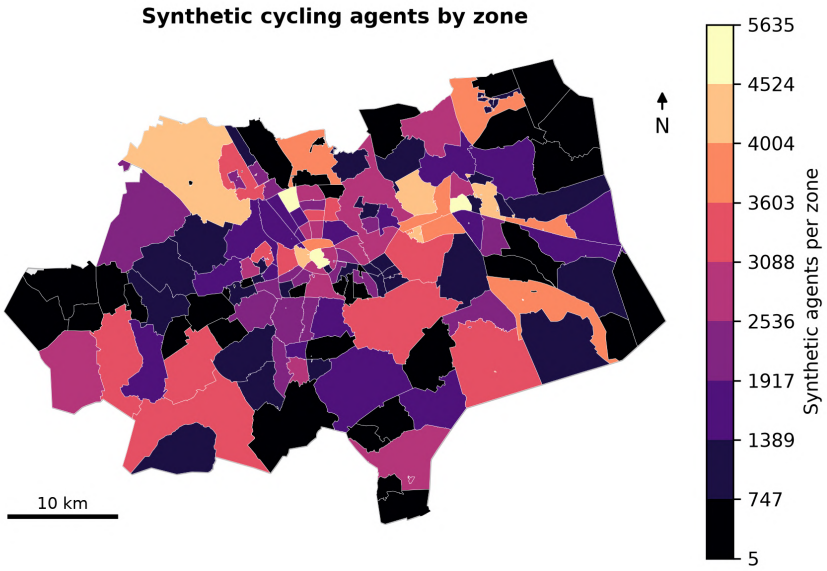

Figure 4.9: Number of synthetic cycling agents per PC4 zone across the MRE.

The two untargeted quantities, the cyclist count and the resulting trips per cyclist, both fall within 1% of the ODiN weights, and the modeled cycling participation rate of 32.1% matches the survey's 31.9%. The synthetic population therefore reproduces not only the demand level it was anchored to, but also the per-person trip rate and the participation rate that it was not, and these are the properties that matter for the assignment that follows.

4.3 Trip Distribution

The regional trip production of Section 4.1 was distributed to destinations by the purpose-specific gravity model of Section 3.4. Each activity purpose carries its own distance-decay parameter β, calibrated so that the modeled cycling trip-length distribution reproduces the observed ODiN distribution for that purpose.

4.3.1 Calibrated Distance Decay

Calibration ran in two stages: a direct fit to the ODiN trip lengths, then a single post-assignment correction. The correction is needed because the network router replaces straight-line distances with routed paths, which are longer. It is largest for the everyday, short-distance purposes, where the work β rose from 0.29 to 0.45 and the shopping β from 0.68 to 1.10, and it is small for touring, whose near-zero β reflects a purpose that is by its nature insensitive to distance. The final parameters are reported in Table 4.5, in which a larger β is a steeper decay and therefore a shorter typical trip.

The effect of these parameters is visible in the assigned trip lengths per purpose (Figure 4.10), where each colored histogram is one purpose, scaled to the same area so that their shapes can be compared: shopping, personal and education trips concentrate within a few kilometers, whereas touring extends far out along the recreational routes.

Purpose	β
Shopping	1.104
Personal	0.962
Other	0.796
Social	0.696
Leisure	0.584
Education	0.568
Work	0.448
Touring	0.067

Table 4.5: Calibrated distance-decay parameter β per activity purpose, after the post-assignment correction.

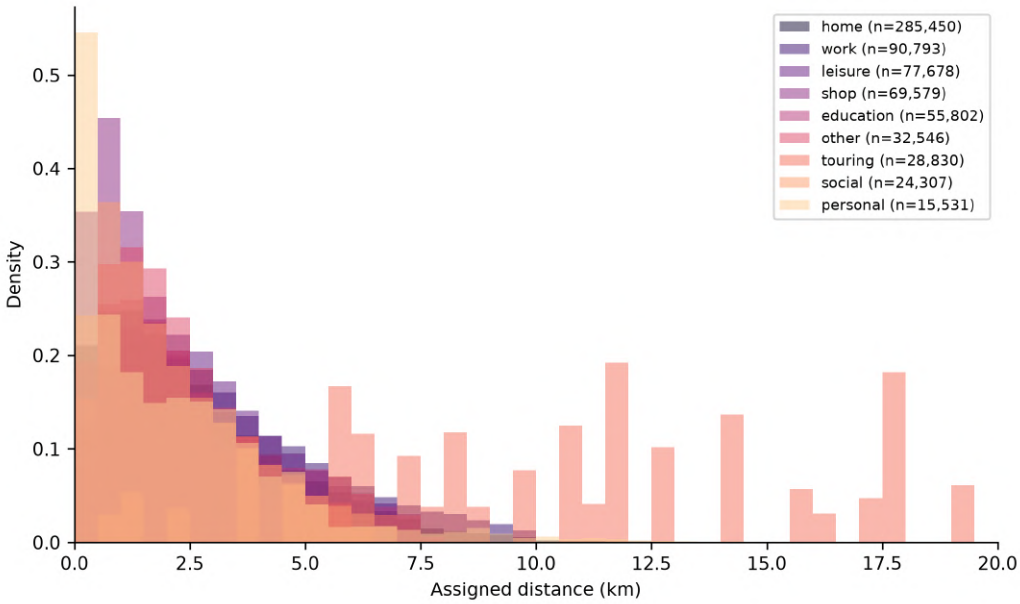

Figure 4.10: Gravity-assigned trip-length distribution per activity purpose.

4.3.2 Assigned versus Survey Distances

Across all purposes together, the gravity-assigned distances reproduce the shape of the ODiN survey distribution, its short-distance peak and long right tail (Figure 4.11). The assigned mean of 3.11 km (straight-line) sits below the survey mean of 4.10 km (as travelled). The gap is what remains after the two-stage calibration, together with the difference between straight-line and travelled distance.

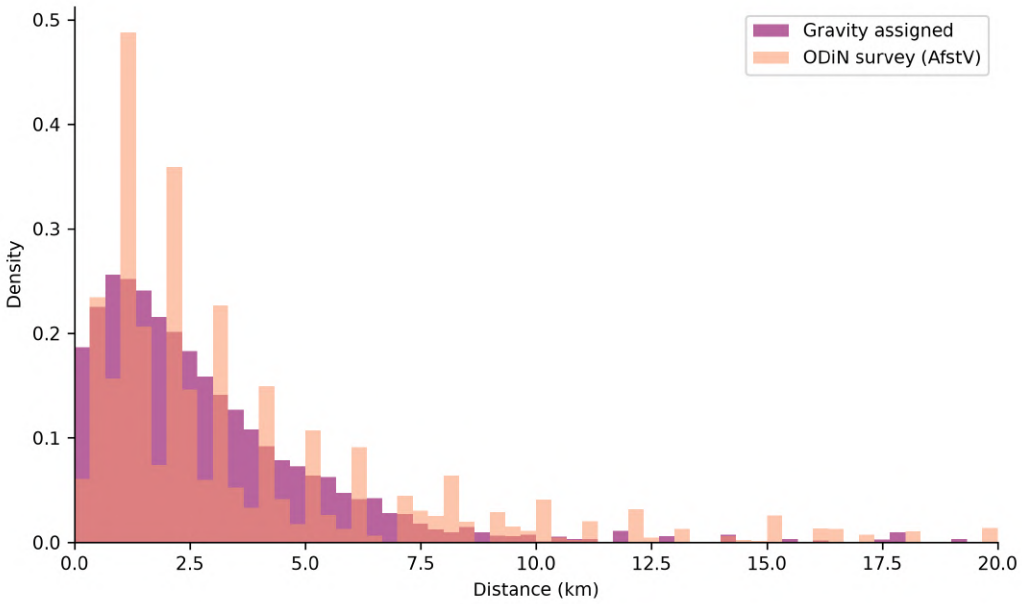

Figure 4.11: Gravity-assigned trip-length distribution against the ODiN survey distances (AfstV).

The composition of the assigned demand by activity purpose, and the assigned distance within each purpose, is given in Table 4.6. Home-based and work trips make up the bulk of the demand, together 55.3% of all assigned trips, and touring stands apart on distance, with a mean assigned length of 14.19 km against 1.56–3.35 km for every other purpose.

Purpose	Trips	Share	Mean (km)	Median (km)
Home	285,450	42.0%	2.71	2.18
Work	90,793	13.3%	3.35	2.82
Leisure	77,678	11.4%	2.87	2.50
Shop	69,579	10.2%	1.56	1.25
Education	55,802	8.2%	2.59	2.12
Other	32,546	4.8%	2.19	1.89
Touring	28,830	4.2%	14.19	11.67
Social	24,307	3.6%	2.15	1.68
Personal	15,531	2.3%	2.17	1.58

Table 4.6: Assigned cycling trips per activity purpose: count, share of the 680,516 assigned trips, and straight-line assigned distance.

4.4 Network Assignment and the Demand Surface

The synthetic population was assigned to the cycling network by the co-evolutionary Multi-Agent Transport Simulation (MATSim) run of Section 3.4. This section reports the network, the convergence of the run, and the link-level cycling-demand surface it produces, the object that Phase 2 intercepts.

4.4.1 Network and Convergence

The routable network is the flat-speed cycling network: 296,729 links together 22,626 km. Every standard ridable link, including all cycleways, residential streets and tracks, carries the same cruising speed of 21.1 km/h, which covers 91.5% of the network length. Lower speeds apply only where cycling is slowed by design: service roads and living streets at 17 km/h (8.0% of the length), footways, pedestrian areas and bridleways at 15 km/h (0.5%), and 36 flights of steps at 3 km/h. Route choice is therefore driven mainly by distance, network shape and the link attractiveness of Equation (3.3), and differences in free-flow speed play a role only on those slower link types. The run used a single mode (bicycle) over 100 iterations at full population, using a reproducible seed of 24.

The run reached a settled state. The average executed plan score rose from 40.5 at iteration 0 to 62.8 at iteration 100, and by the last iterations the executed score sits within 0.2 of the best-known score (Figure 4.12, where the three lines give the average executed, best and worst plan score per iteration), so the replanning has stopped finding clearly better plans. Of the planned trips, 94.9% are executed, and the remaining 5.1% are trips the network cannot clear, an effect of the storage-capacity floor on this unusually short-linked network, which is taken up as a limitation.

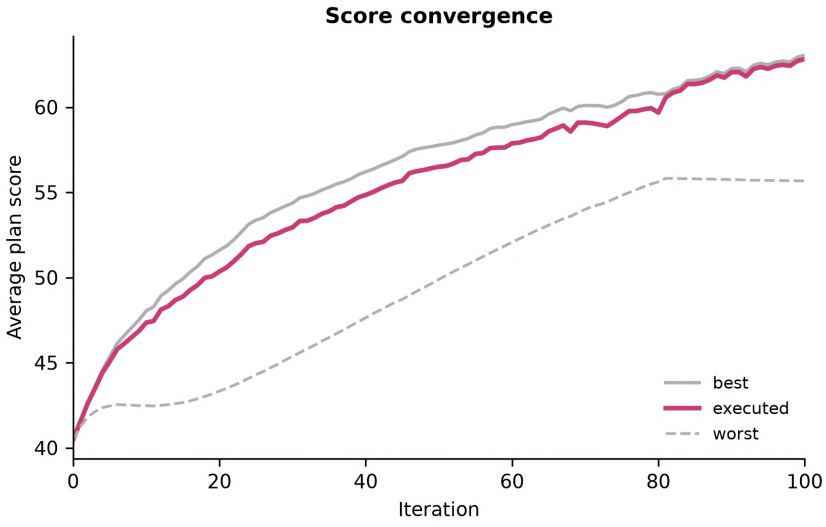

Figure 4.12: Average plan score (executed, best, worst) over the 100 replanning iterations.

4.4.2 The Link-Level Demand Surface

The assignment loads 237,974 of the network's links (80.2%), together 17,251 km, with 44.6 million link crossings on an average weekday. The load is very uneven: the median loaded link carries only 38 crossings a day, while the busiest carries 7,010 and the 99th percentile (2,041) is more than fifty times the median (Table 4.7). Ranking the links from

busiest to quietest shows the same thing directly: the busiest 7% of loaded links carry half of all crossings, and the busiest 20% carry 80% (Figure 4.13). Demand therefore falls on a thin set of heavily used corridors against a diffuse background. This is the spatial structure that makes the flow-capturing siting of Phase 2 useful, because a small number of well-placed points can intercept a large share of all movement.

Quantity	Value
Network links / length	296,729 / 22,626 km
Loaded links / length	237,974 (80.2%) / 17,251 km
Total link crossings / weekday	44,622,909
Median loaded-link volume	38
90th / 99th percentile	516 / 2,041
Maximum link volume	7,010

Table 4.7: Distribution of average-weekday cycling volume over loaded network links, final full-population run.

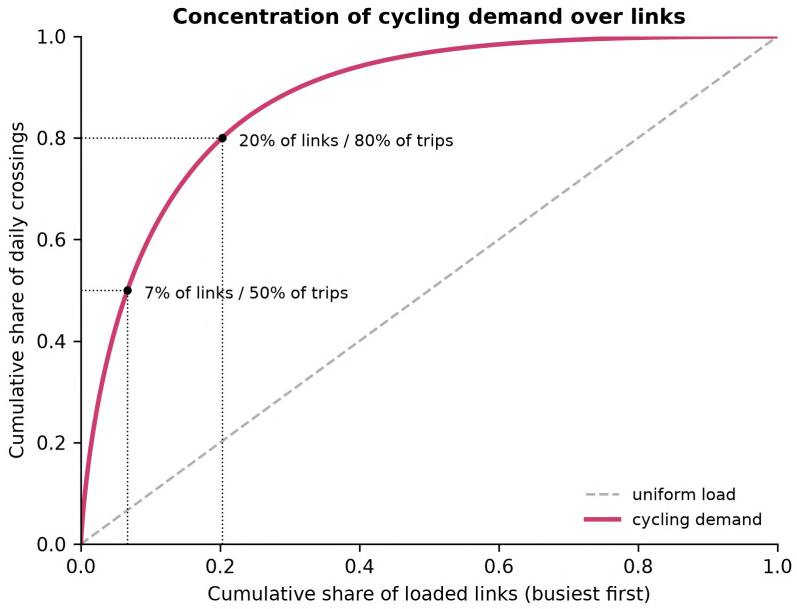

Figure 4.13: Cumulative share of average-weekday crossings against the cumulative share of loaded links, ordered busiest first. The dashed line is a uniform load.

The same demand mapped onto the network is shown in Figure 4.14. Average-weekday cycling volume concentrates on a connected set of main corridors through the urban cores and thins to a diffuse background over the rest of the network, the spatial form of the concentration measured in Figure 4.13.

The busiest 1% of loaded links, the corridors carrying the heaviest cycling, are mapped in Figure 4.15. They form a small, connected skeleton concentrated in the Eindhoven core and along the radial routes into it, with a second cluster in Helmond, against the far larger background of ordinary links that carry little traffic. This is the demand field of Figure 4.14 reduced to its busiest corridors alone.

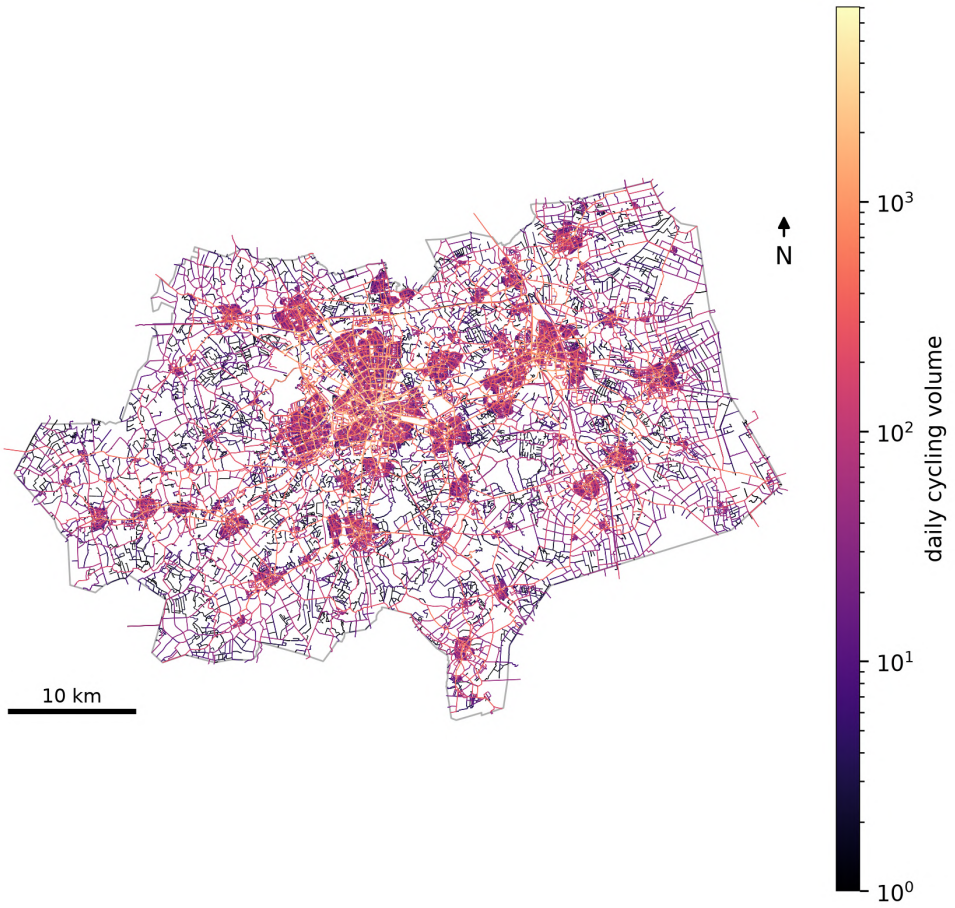

Figure 4.14: Average-weekday cycling demand across the MRE network, links coloured by daily volume (log scale).

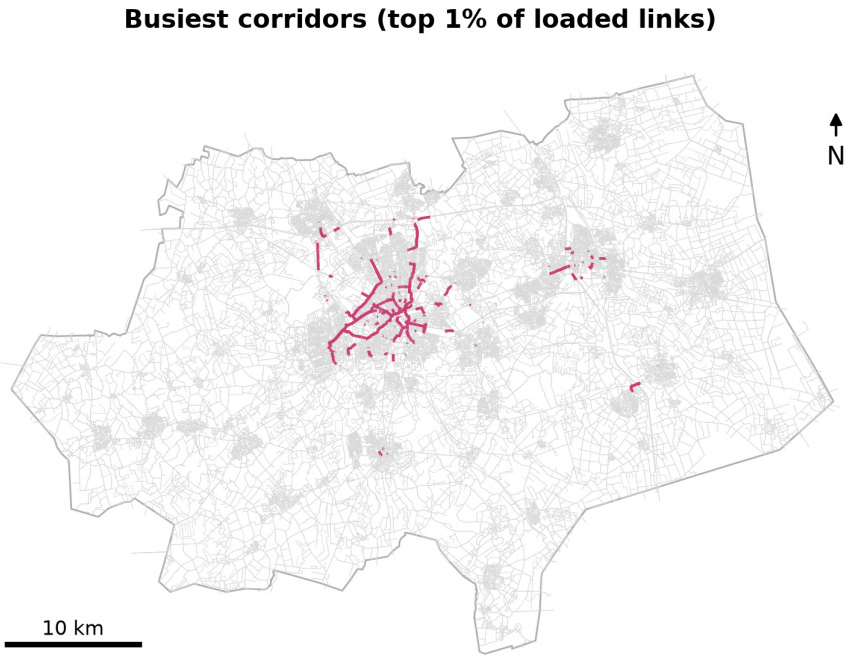

Figure 4.15: The busiest 1% of loaded links (those carrying at least 2,041 crossings a day) against the rest of the loaded network.

4.4.3 Temporal Profile

The assigned demand shows a sharp morning commuting peak at 08:00 and a broad band of afternoon activity that is highest around 14:00 before easing through the evening (Figure 4.16). In the figure each bar is one hour of departures, the dark front bars are the simulated legs and the lighter bars behind are the observed counts, and pink marks the two-hour commute windows 07:00–09:00 and 16:00–18:00. These are the same windows the Phase 2 commuter model intercepts (Section 3.5.2), and together they account for 31.0% of executed legs and 30.6% of link volume. One feature runs against observation: modeled trip length rises a little through the day (from 3.96 km in the morning window to 4.62 km in the evening) where ODiN trips shorten, because the gravity model assigns distance without regard to departure hour, and this is taken up in the validation.

The observed Nationaal Dataportaal Wegverkeer (NDW) profile, drawn behind the simulated bars in Figure 4.16 and scaled to the same daily total, shows where the simulated timing departs from observation. The morning peak falls in the right hour, 08:00, but is lower than observed: the 07:00–09:00 window holds 17.5% of the simulated daily total against 20.6% of the counts. The larger difference is in the afternoon. The counts rise to a second peak at 16:00 and stay high at 17:00, while the simulated afternoon peaks two hours earlier, at 14:00, and the 16:00–18:00 window holds 13.6% of simulated legs against 18.3% observed. The volume missing from the late afternoon sits in the late morning and early afternoon instead, where 09:00 to 15:00 holds 48.5% of simulated legs against 40.7% observed.

This early afternoon peak is inherited from the survey diaries rather than produced by the simulation. The planned departures of the synthetic population already peak at 14:00 (7.9% of trips) before any replanning, because each agent copies the timing of a real ODiN respondent. The respondents travelling at that hour are a different group from those at the commuter peak: at 14:00, 28% of planned trips are made by respondents under 18,

largely the return from school, and 26% by respondents aged 65 or over, while at 17:00 the working-age share is 65% against 46% at 14:00. The counting stations, most of which sit on routes between towns, are expected to be weighted toward that working-age commute and therefore peak later. The simulated profile describes all weekday cycling in the region, and the counts describe cycling on the main inter-town routes, so the two-hour offset in the afternoon peak reflects the mix of trips in the diaries rather than a timing error introduced by the assignment.

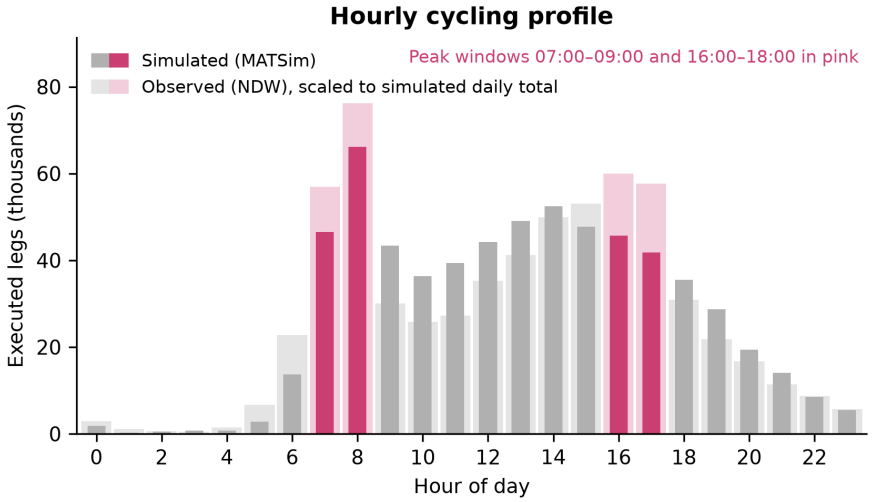

Figure 4.16: Hourly distribution of executed cycling legs over an average weekday (front bars), with the hourly NDW counts over the 47 counting stations scaled to the same daily total (lighter bars behind). The commute windows 07:00–09:00 and 16:00–18:00 are shown in pink.

The lengths of the executed trips are summarised in Figure 4.17. The distribution is short and right-skewed: the number of trips rises to a peak between about 1 and 3 km and then falls away, leaving a long tail of occasional longer trips. The median executed trip is 3.11 km and the mean is 4.17 km. The mean sits above the median because the long tail pulls it upward, which is the expected shape for urban cycling.

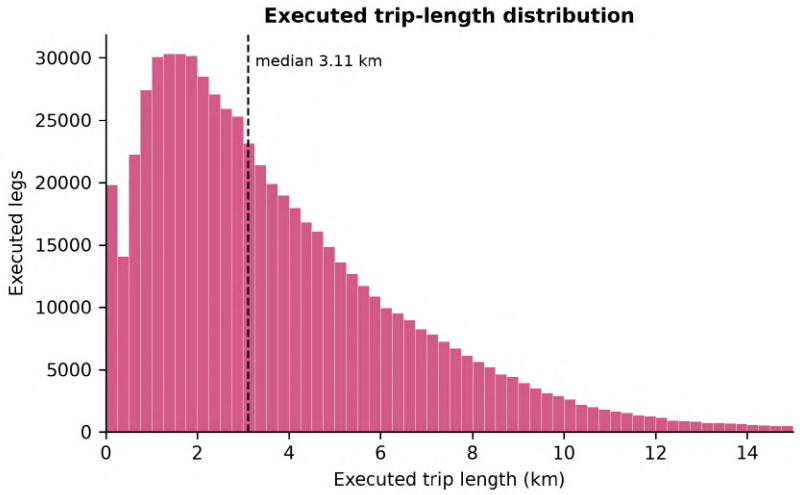

Figure 4.17: Distribution of executed cycling trip lengths over the simulated network, with the median (3.11 km) marked. Trips longer than 15 km are omitted from the horizontal axis for legibility.

4.5 Model Validation

The assigned demand surface is validated against the 47 NDW counting stations of Section 3.4, giving $47 \times 24 = 1{,}128$ station-hour comparisons. All figures are computed on the corrected one-link-per-direction cross-section.

4.5.1 Overall Fit

Table 4.8 reports the fit of the final run, and Figure 4.18 plots each simulated station-hour against its observed count. The model reproduces the hourly counting-station volumes with a Nash–Sutcliffe efficiency of 0.518, a Pearson correlation of 0.774 ($r^2 = 0.60$) and a volume ratio close to one (Sim/Obs 1.067). By Geoffrey E. Havers (GEH) (3.18), the usual measure for traffic-count validation, 78.1% of station-hours fall below 5 and the mean GEH is 3.38. In the scatter the cloud hugs the 1:1 line, tightest at the low and mid volumes where most station-hours lie, with the GEH<5 points forming that diagonal core and the poorer matches scattering above and below it.

Against the Design Manual for Roads and Bridges (DMRB) standard of GEH<5 for at least 85% of counts (Highways Agency, 1996), the model does not pass, reaching 78.1%. The looser threshold of GEH<10, above which a comparison is taken to indicate a probable mismatch, is met by 94.3% of station-hours, so most of the misses are moderate rather than gross. The other measures place the fit in the acceptable range. A mean SQV of 0.905 falls in the highest of the quality classes used for the SQV in agent-based model validation, which rate values of 0.9 and above as excellent and values below 0.7 as unsatisfactory (Scherr et al., 2020). The NSE of 0.518 lies above the value of 0.5 that the hydrological literature, where the measure originates, treats as satisfactory (Moriasi et al., 2007). The fit is therefore reported as acceptable for a strategic, region-scale cycling model rather than as meeting DMRB.

Figure 4.19 shows the distribution of GEH over the station-hours. Most of the mass sits below 5, with a long tail of a few poorly-matched station-hours, and the 5 and 10 reference thresholds are marked. Figure 4.20 compares the mean observed (grey) and simulated (pink) volume per station by hour of day. The observed series has the clearer two-peak commuting shape, with peaks at 08:00 and 16:00, while the simulated series peaks in the morning at 08:00 but places its second peak earlier, around 14:00. The simulated volume

Metric (hourly)	Value
Nash–Sutcliffe efficiency (Nash–Sutcliffe efficiency (NSE))	0.518
Coefficient of determination r^2	0.599
Pearson r	0.774
Spearman ρ	0.858
GEH<5	78.1%
GEH<10	94.3%
Mean GEH	3.38
Mean Scalable Quality Value (SQV)	0.905
Volume ratio (Sim/Obs)	1.067
Root Mean Squared Error (RMSE) / Mean Absolute Error (MAE)	44.1 / 25.4
Trips executed	94.9%

Table 4.8: Fit of the final full-population run against the 47 NDW counting stations, hourly ($n = 1{,}128$ station-hours).

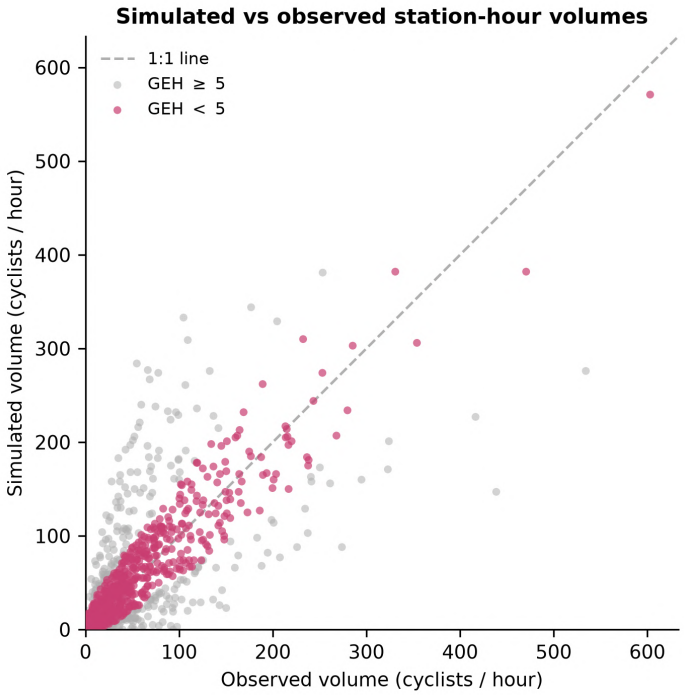

Figure 4.18: Simulated against observed cycling volume for each of the 1,128 station-hours, with the 1:1 line. Points with GEH<5 are shown in colour.

sits below the observed at both the morning and the late-afternoon peaks and runs above it through the early afternoon.

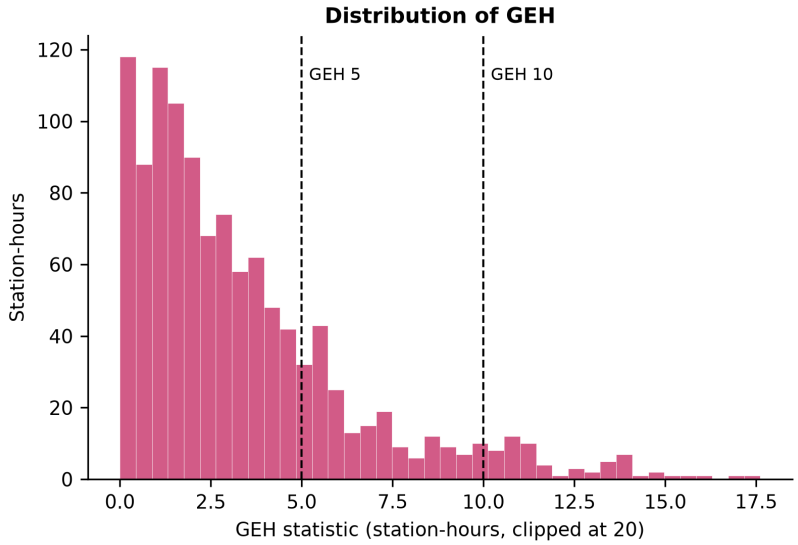

Figure 4.19: Distribution of the GEH statistic over the 1,128 station-hours, clipped at 20. The 5 and 10 thresholds are marked.

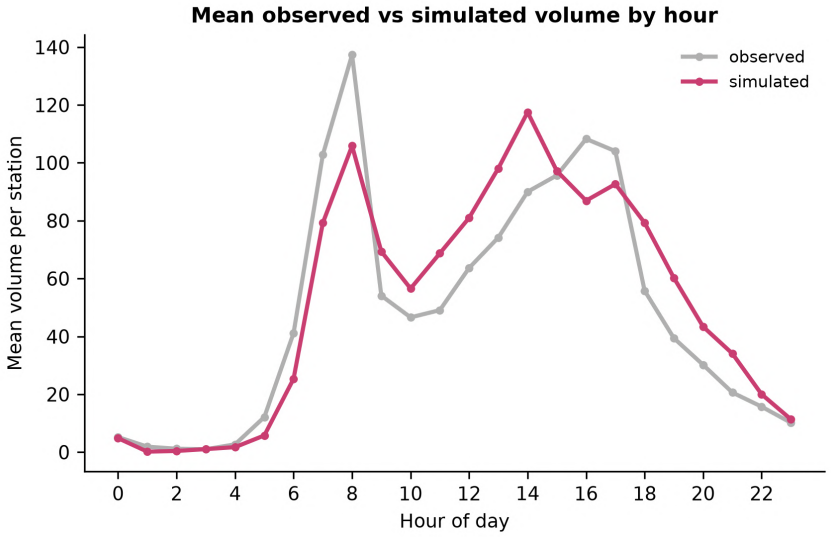

Figure 4.20: Mean observed and simulated volume per counting station, by hour of day.

4.6 Conclusion

This chapter established the following, in the order the model produces them.

Trip generation. A negative binomial model estimated on 113 observed PC4 zones retains two covariates, land-use mix and cycling-lane density, each raising the expected per-capita cycling rate by 21% and 15% respectively for a one-standard-deviation increase. Against a

population-only model the two add $\chi^2(2) = 10.6$ ($p = 0.005$), and out-of-sample R^2 across six folds ranges from -0.80 to $+0.57$. Anchoring the fitted rates to the ODiN control total moves the regional total from 579,082 to 680,139 trips, or 0.86 bicycle trips per inhabitant on an average Monday-to-Friday day.

Population synthesis. Donor matching produces 254,010 cycling agents carrying 679,448 bicycle trips at full population scale, with a cycling participation rate of 32.1% against the survey's 31.9%.

Trip distribution. The purpose-specific gravity model reproduces the shape of the survey distance distribution, at an assigned straight-line mean of 3.11 km against a surveyed travelled mean of 4.10 km.

Trip assignment. The routable network holds 296,729 links over 22,626 km, of which the assignment loads 237,974 links (80.2%) over 17,251 km with 44.6 million link crossings. Load is concentrated: the busiest 20% of loaded links carry 80% of all crossings. The mean executed plan score rises from 40.5 to 62.8 and settles inside the 2% convergence criterion, with 94.9% of planned trips executed. The morning and evening peak windows carry 31.0% of executed legs and 30.6% of link volume.

Validation. Against 1,128 station-hours from NDW counters the model returns a Nash–Sutcliffe efficiency of 0.518, $r^2 = 0.599$, a volume ratio of 1.067, a mean GEH of 3.38 and a mean SQV of 0.905. It reaches GEH<5 for 78.1% of station-hours, below the DMRB threshold of 85%, and GEH<10 for 94.3%.

The output carried into Chapter 5 is the hourly, link-level demand surface these steps produce, together with the reconstructed path of every simulated trip.

Chapter 5

Phase 2 Results: Siting of Cycle Service Points

This chapter applies Phase 2 of the framework to the demand surface produced in Chapter 4. It follows the model's own order: the reconstructed demand and the two demand models it is run under, the candidate set and its suitability surface, the eight quality factors and their overlap, the coverage attainable as a function of the number of Cycle Service Points (CSPs) built, the structure of the selected sites and their agreement across models, and finally how far that result survives future demand growth and the parameters the model fixes by choice. The recommendation is a ranked set of corridor locations rather than exact coordinates to build, and the sites are reported at that resolution.

5.1 The Demand Object and the Two Models

Phase 2 sites cycle service points to intercept the cycling flow produced in Phase 1. The model is a Flow-Capturing Location Model (FCLM): sites are chosen to maximise the realized cycling flow they capture, and a trip captured at a higher-quality position is worth more than the same trip at a poor one. The objective for a set of sites Y of size p is

$$Z(Y) = \sum_{\text{trips } t} q_t \max_{j \in Y} \big[\kappa(d_{tj}) Q_j\big], \tag{5.1}$$

where $\kappa(d)$ is a capture kernel decaying linearly to zero at 150 m, and $Q_j \in [0,1]$ is the quality of position j. Each trip is counted once, by its single best site, so two sites on one corridor compete for the same flow and anti-clustering is a property of the objective rather than a penalty.

The demand is reconstructed from the frozen Multi-Agent Transport Simulation (MATSim) event stream of Phase 1 as realized link-by-link paths, one path per leg, so that the model knows whether two links carry the same cyclist. Two demand windows are run (Table 5.1): a `commuter` model on the peak trips and a `baseline` model on all-day trips. Flow is only ever the objective, and the daily/peak distinction is expressed by the demand window, not by adding a flow factor, which preserves the flow-capture contribution of the method.

The commuter window carries $195{,}513$ peak trips against the baseline's $626{,}944$ all-day trips, and the peak flow the commuter model intercepts is mapped in Figure 5.1, where brighter links carry more cyclists. It is concentrated on the corridors through the Eindhoven core and the secondary town centers.

The all-day flow that the baseline model intercepts is shown in Figure 5.2. It carries the same corridor structure as the peak field of Figure 5.1, but the flow is less concentrated: all-day cycling spreads onto the residential and shopping streets that stay faint in the

	Commuter (peak)	Baseline (all-day)
Demand trips captured	195,513	626,944
Candidate positions	235,844	269,379

Table 5.1: The two demand models: intercepted flow and candidate positions.



Figure 5.1: Reconstructed peak cycling flow over the Metropolitan Region Eindhoven (MRE) network, links coloured by volume (log scale).

peak field, so the two models see a similar skeleton of main corridors against a differently weighted background.

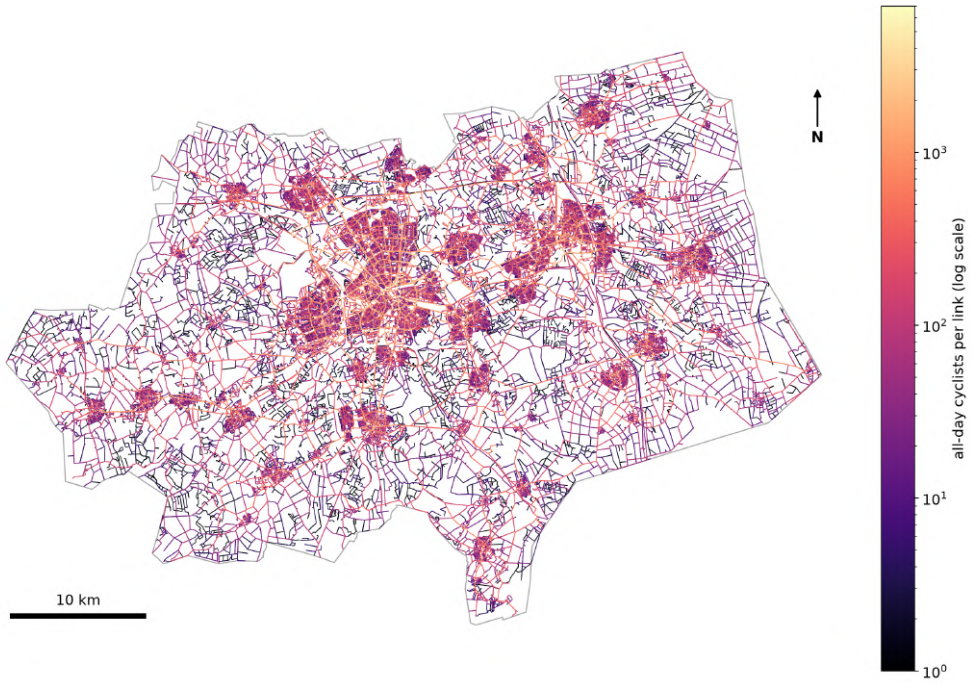

Figure 5.2: Reconstructed all-day cycling flow over the MRE network, links coloured by volume (log scale).

5.2 Candidates and the Capturable-Flow Surface

Every link-adjacent position that a service point could occupy is a candidate: 235,844 for the commuter model and 269,379 for the baseline model. Between them these positions reach every trip in their demand window, so no demand is out of range of the candidate set and the coverage figures below are limited by the budget p rather than by where a service point may stand. Before any quality weighting, each candidate can be scored on the share of all-day cycling flow that a single service point there would capture on its own. Figure 5.3 maps this capturable-flow share, where brighter colors mean a larger share. It is the flow side of Equation (5.1): the quality score Q_j of Section 5.3 then scales it but never overrides it. Capturable flow is highest in a bright core over Eindhoven and along the corridors radiating from it, with secondary hotspots in Helmond and the other town centers, and it falls away across the rural network. The best single position captures 2.06% of all-day flow on its own.

The demand reconstruction and the candidate generation each pass through a transparent filter (Table 5.2). On the demand side almost all reconstructed legs survive: only trips shorter than 300 m are dropped, leaving 626,944 baseline trips. On the candidate side the 296,093 directed links inside the study area collapse to undirected alignments, then to the alignments actually used by at least one trip, and are finally re-sampled at 25 m spacing, giving 269,379 candidate positions for the baseline model and 235,844 for the peak-flow commuter model.

The hotspot is a property of the region rather than of the display grid: the same central concentration is visible in Figure 5.4 at all four display resolutions, from the coarse 1,000 m grid down to 25 m, sharpening but not shifting as the grid refines.



Figure 5.3: Share of all-day cycling flow that a single service point would capture on its own, one point per candidate position.

Stage	Count
Reconstructed legs	644,904
Paths ≥ 300 m (retained trips)	626,944
Directed links	296,093
Undirected alignments	157,155
Links with ≥ 1 passing trip	135,696
Points placed at 25 m	374,491
Thinned to ≥ 25 m (baseline)	269,379
Peak-flow subset (commuter)	235,844

Table 5.2: Trips retained from the MATSim event stream, and candidate positions retained through the generation filter. The candidate stages are for the baseline model, while the commuter model runs the same filter on the peak-flow subset.

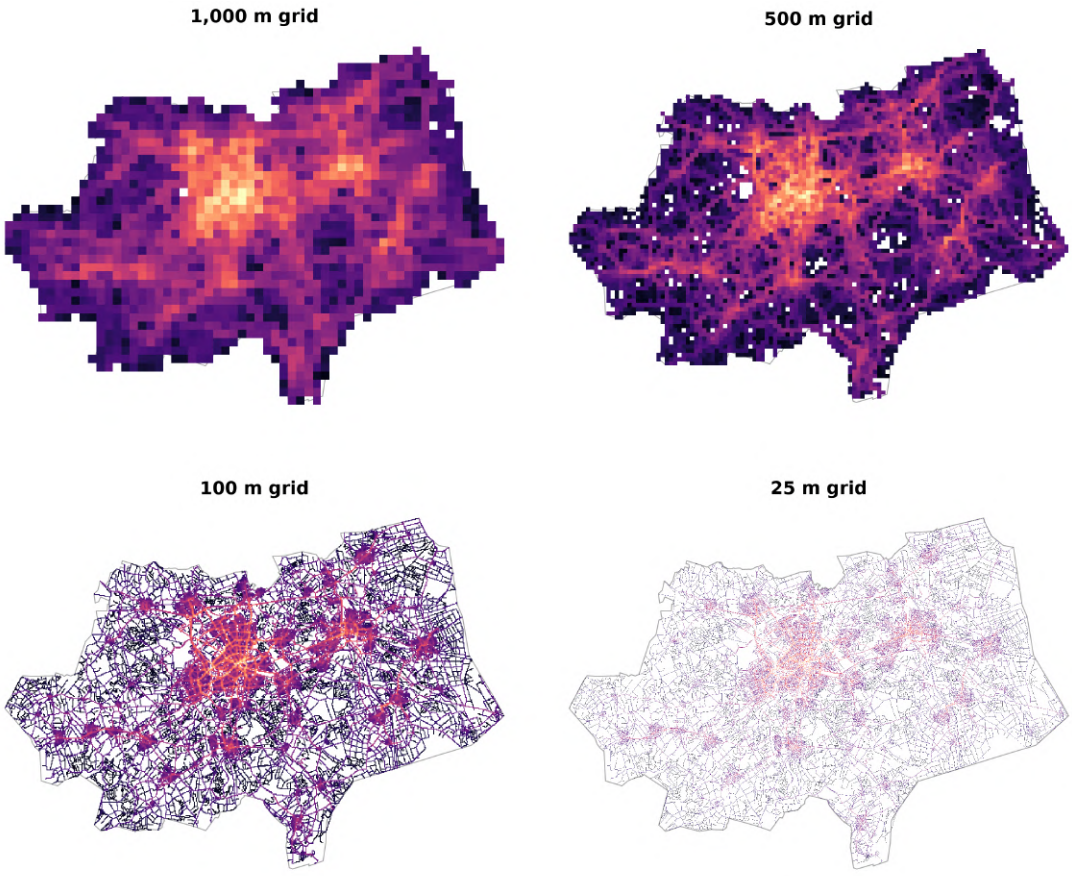

Figure 5.4: The capturable-flow surface at four display resolutions, from 1,000 m to 25 m. Each cell carries the share its best candidate would capture, on one scale shared by all four panels.

5.3 The Eight Factors and Their Overlap

Quality is built from eight factors, each computed on real data over its own spatial reach (Table 5.3). All eight are oriented so that a higher score is more suitable. Urban core and main cycle-network proximity are built from cost-type raw inputs (an urbanity class and a distance to a route) but are scored so that dense, central and on-route positions rank highest.

Key	Factor	Reach	Direction
X	Route-junction proximity	150 m	benefit
C	Betweenness centrality	link (7.5 km cutoff)	benefit
M	Land-use mix	500 m	benefit
U	Urban core (CBS urbanity)	CBS buurt	benefit
B	Business density	1,000 m	benefit
A	Accessibility to opportunities	5,000 m	benefit
H	Transit-hub proximity	1,500 m	benefit
F	Main cycle-network proximity	100 m	benefit

Table 5.3: The eight quality factors, their spatial reach and direction. All are normalized to $[0, 1]$ before weighting.

The spread of each normalized factor across the candidate positions is given in Table 5.4. Most factors use close to the full $[0, 1]$ range with a mean near 0.5, so they separate candidates well. Two are concentrated near zero, urban core (mean 0.21) and main cycle-network proximity (mean 0.25): most candidates score low on these and only a minority score high, so within the weighted quality they mainly reward the few candidates that sit inside the urban core or directly on a main cycle route.

Key	Factor	Mean	SD	Min	Max
X	Route-junction proximity	0.500	0.288	0.077	1.000
C	Betweenness centrality	0.500	0.289	0.000	1.000
M	Land-use mix	0.478	0.251	0.000	1.000
U	Urban core	0.206	0.303	0.000	1.000
B	Business density	0.500	0.288	0.049	1.000
A	Accessibility to opportunities	0.500	0.289	0.000	1.000
H	Transit-hub proximity	0.500	0.288	0.080	1.000
F	Main cycle-network proximity	0.245	0.381	0.000	1.000

Table 5.4: Distribution of each normalized quality factor ($[0, 1]$) across the commuter-model candidate positions.

5.3.1 Factor Overlap

The eight factors are not independent (Figure 5.5). The heaviest overlap is between the urban-core and transit-hub factors (Spearman $\rho = 0.78$), just under the 0.80 redundancy flag, and business density sits almost as close to transit hub (0.77). Seven of the eight factors, route junction, betweenness, land-use mix, urban core, business density, accessibility and transit hub, inter-correlate between roughly 0.24 and 0.78, and each of them rises toward the dense parts of the region. Only main cycle-network proximity stands apart: it correlates below 0.2 in absolute value with all seven others. The eight criteria therefore carry close to

two independent dimensions rather than eight, one of them urban intensity and the other the designated cycle network.

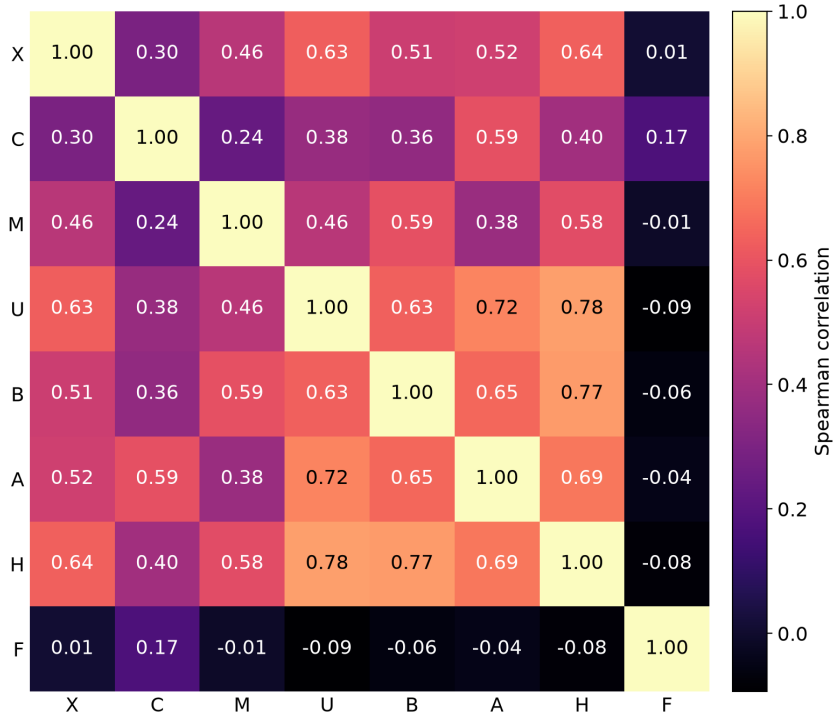

Figure 5.5: Spearman rank correlation between the eight quality factors.

Figure 5.6 maps the eight factors side by side. Its final panel is the capturable-flow surface, the share of cycling flow the best site in each cell would intercept, which is the model's objective and is shown here for comparison with the factors that modulate it. The panels carry the same split as the correlation matrix. Route junction, betweenness, land-use mix, urban core, business density, accessibility and transit-hub proximity all brighten over the central urban area and over the secondary town centers, while main cycle-network proximity picks out a thin linear pattern that follows the designated routes between towns and ignores the urban cores. Accessibility is the widest and smoothest of the seven, because its 5,000 m reach and slow distance decay blur the individual town centers into one regional gradient, but its peak sits over Eindhoven with the rest.

5.3.2 Where Each Factor Would Site

To show what each factor contributes, the siting was re-run using each factor's own spatial pattern alone, ignoring flow, and the captured flow compared against demand-only siting. The result is the clearest evidence for the flow-first design (Table 5.5). Siting on demand alone captures 9.2% of flow at $p = 5$, while siting on any single factor's geography captures at most 3.5%. Betweenness comes closest, at 3.46%, because topological centrality partly coincides with where cycling concentrates, and accessibility follows at 2.55% because its regional gradient peaks over the same urban core. Land-use mix collapses to 0.30%: the places that score highest on it are villages and small town centers rather than the corridors that carry the cycling. No quality factor is a substitute for the flow objective. The factors refine which of several high-flow positions is chosen. They do not decide where the flow is.

Figure 5.7–5.14 show, for each factor, its surface (left) and where the p sites would fall if only that factor mattered (right, numbered pink points) against the demand-only sites (open

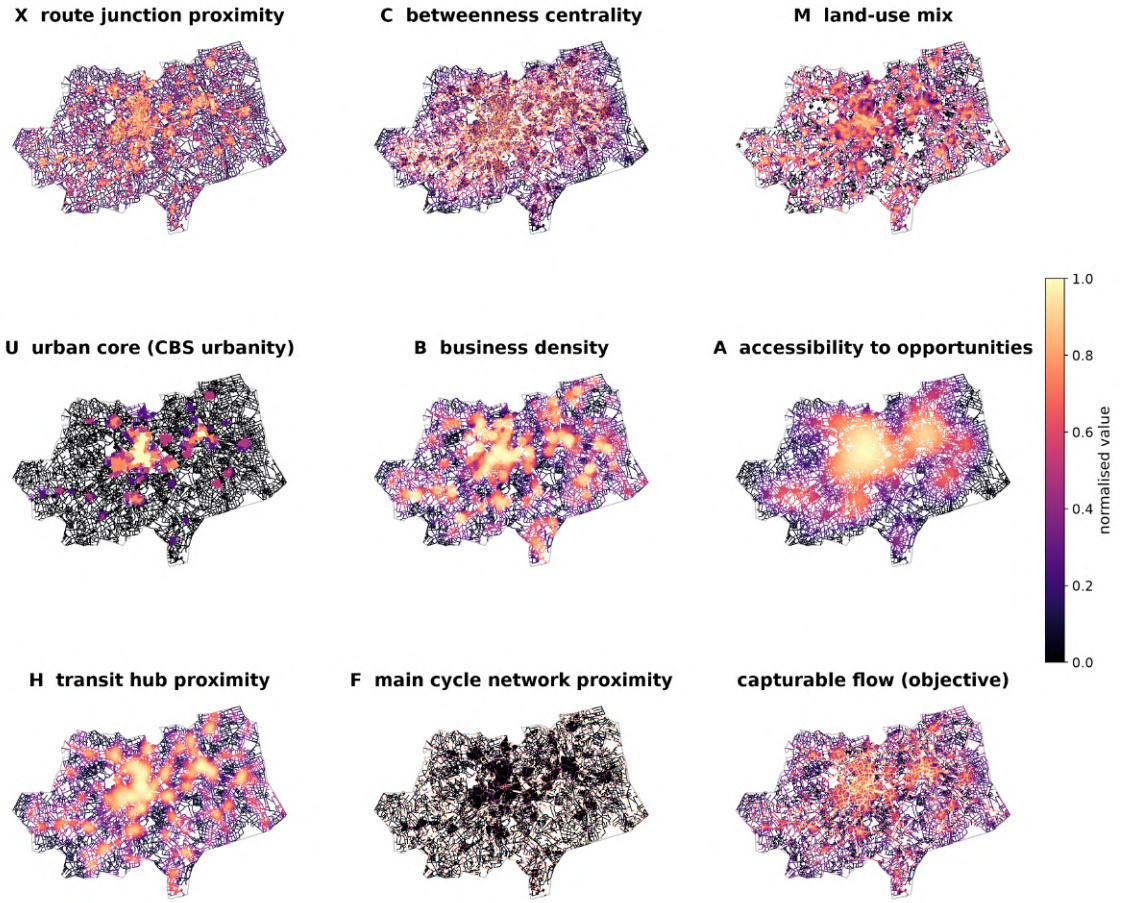

Figure 5.6: The eight quality factors across the region, one panel each, at candidate resolution. The ninth panel is the capturable-flow objective surface.

	Factor	Flow captured
–	Demand only (flow objective)	9.21%
C	Betweenness centrality	3.46%
A	Accessibility to opportunities	2.55%
H	Transit-hub proximity	1.46%
X	Route-junction proximity	1.39%
B	Business density	0.89%
U	Urban core	0.83%
F	Main cycle-network proximity	0.45%
M	Land-use mix	0.30%

Table 5.5: Flow captured at $p = 5$ when siting by each factor's spatial pattern alone, and by demand alone (commuter model).

black circles). Because the single-factor sites are forced at least 10 km apart (Section 3.5.5), the right-hand panels show how a factor ranks the region as a whole rather than where its highest values cluster.

Betweenness (Figure 5.7) illustrates this. Its surface peaks in and around Eindhoven, so the first single-factor site is the highest-scoring position in the city. Once it is taken, the 10 km separation excludes the rest of Eindhoven, so the remaining four sites go to the best positions outside that radius, near Helmond, Son, Bladel and Someren. These still rank high among all candidates, because betweenness also rewards the links that many shortest routes between towns must pass. The spread of these sites over the rural part of the region is therefore a result of the separation rule, not a preference of betweenness for rural locations.

The other factors read the same way. Accessibility (Figure 5.8), transit-hub proximity (Figure 5.10), urban core (Figure 5.12) and business density (Figure 5.11) all put their first two sites in Eindhoven and Helmond, the same two municipalities the demand-only solution uses, because opportunities, stations, dense fabric and employers all concentrate there. They then differ in where the separation rule sends the remaining three. Accessibility and urban core keep to town centers, in Valkenswaard, Best, Asten, Someren and Gemert, while business density falls back on industrial estates such as those in Helmond and Bladel, which are dense in employers but thin in cycling. Route-junction proximity (Figure 5.9) is high across the built-up area of every town, so it places no site in Eindhoven at all and spreads across Helmond, Veldhoven, Geldrop, Bladel and Son. Land-use mix (Figure 5.13) scores village centers about as high as the city and sends all five sites to Gemert, Bergeijk, Veldhoven, Cranendonck and Asten, and main cycle-network proximity (Figure 5.14) is high along the designated routes between towns rather than inside the built-up areas, so all five of its sites fall on scattered housing outside any town. These last two capture almost no flow.

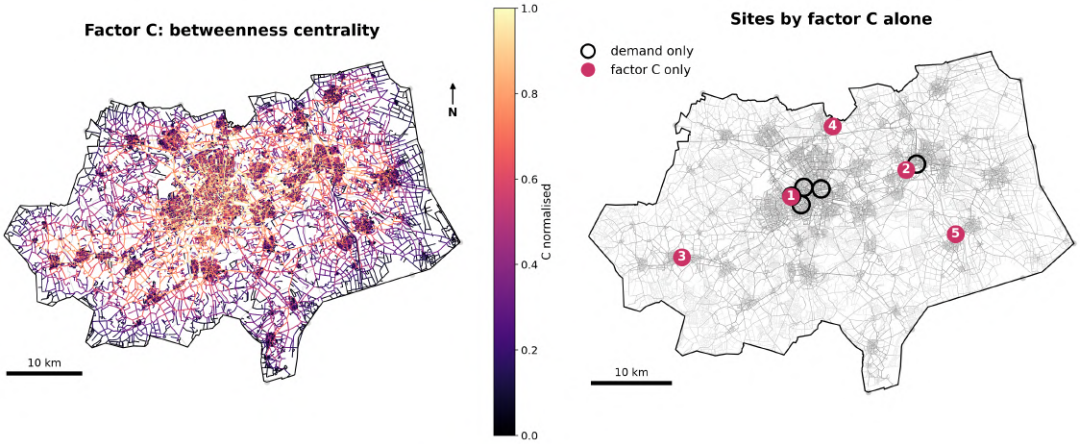

Figure 5.7: Betweenness centrality (C): surface and single-factor siting.

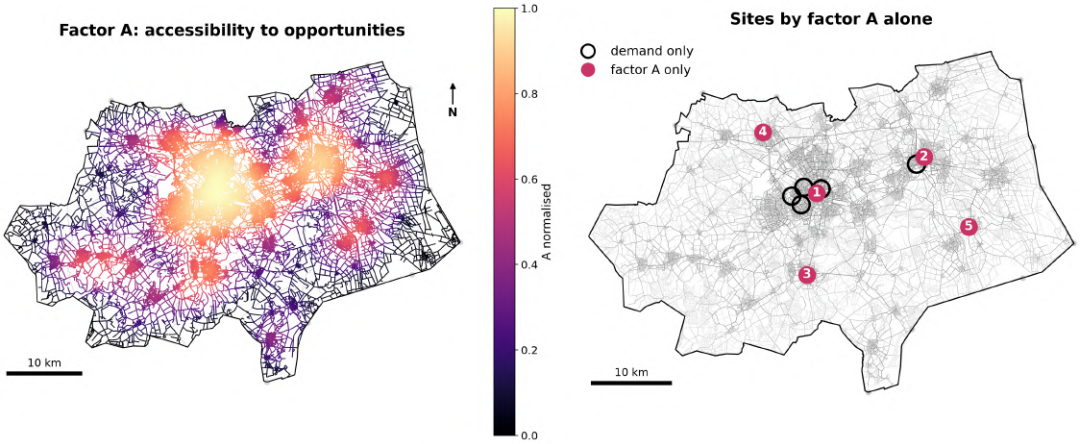

Figure 5.8: Accessibility to opportunities (A): surface and single-factor siting.

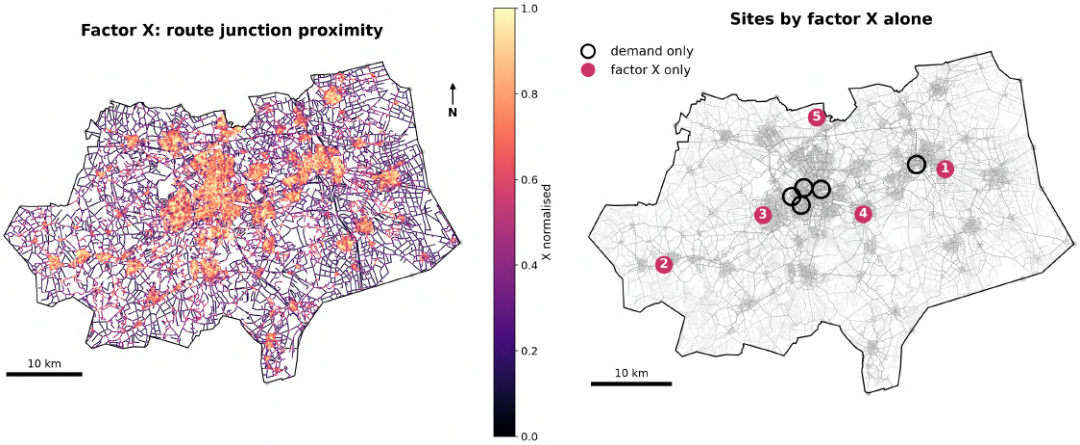

Figure 5.9: Route-junction proximity (X): surface and single-factor siting.

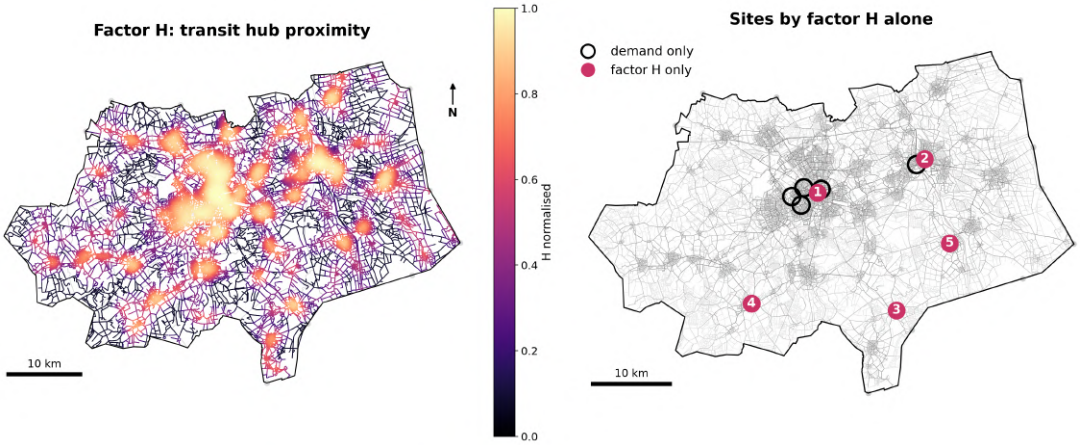

Figure 5.10: Transit-hub proximity (H): surface and single-factor siting.

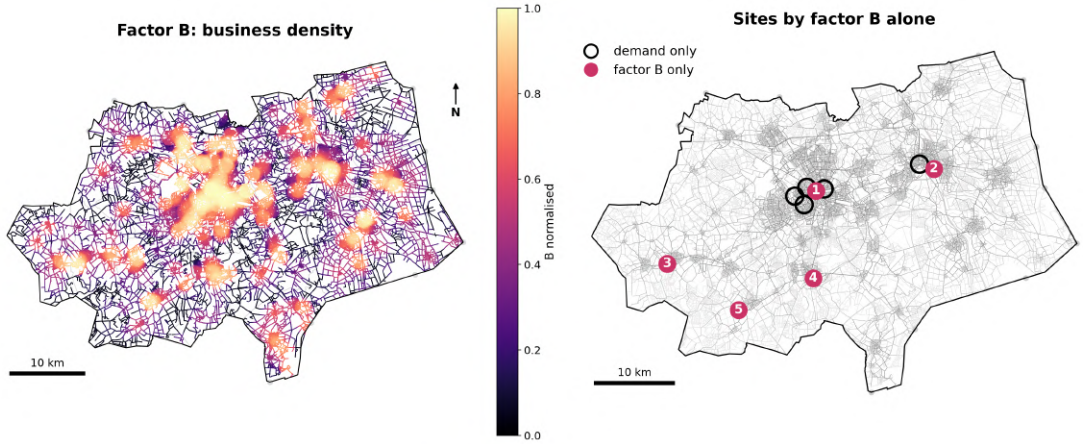

Figure 5.11: Business density (B): surface and single-factor siting.

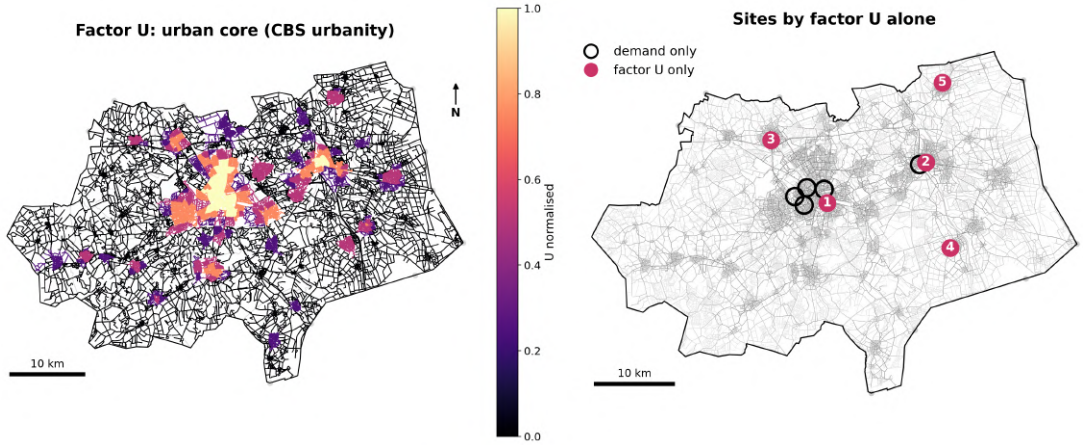

Figure 5.12: Urban core (U): surface and single-factor siting.

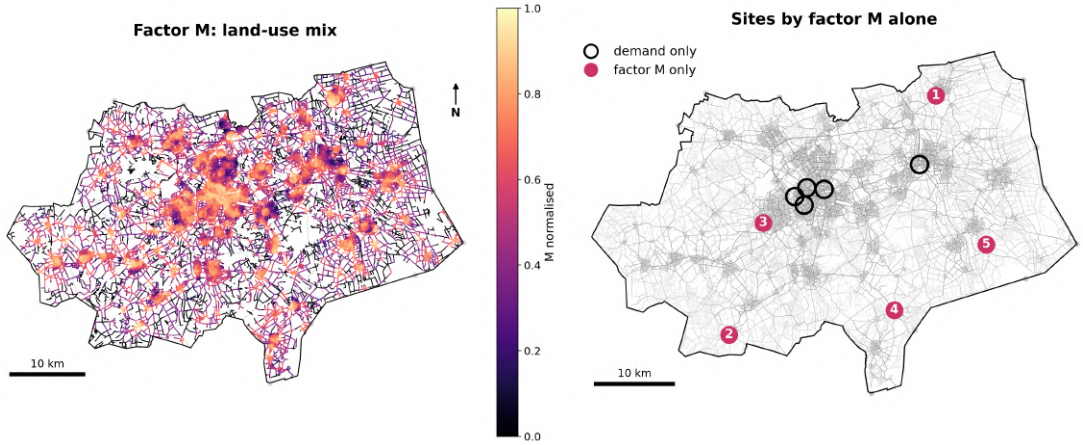

Figure 5.13: Land-use mix (M): surface and single-factor siting.

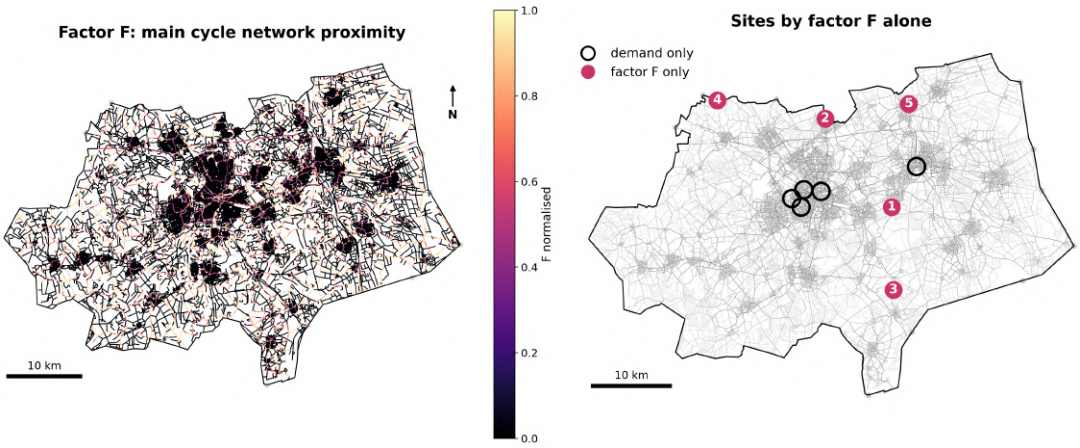

Figure 5.14: Main cycle-network proximity (F): surface and single-factor siting.

5.4 Coverage

Table 5.6 reports the share of demand captured by the optimized sites for each model at four budgets p. The commuter model captures more than the baseline at every p, by one percentage point at $p{=}3$ and $p{=}5$ and by 1.6 points at $p{=}10$: peak flow concentrates on commuter corridors, so fewer sites intercept a larger share, while all-day flow is more diffuse.

Five sites capture 8.5% of peak flow in the commuter model and 7.6% of all-day flow in the baseline model, about 16,700 of the 195,513 peak trips and 47,000 of the 626,944 all-day trips on an average weekday. For five point facilities in a network of 22,000 km this is a substantial share, and it follows from the concentration of demand on a small number of corridors (Section 4.4.2). The coverage curve rises close to linearly from three sites to ten, with each added site capturing between 1.0 and 1.3 percentage points of peak flow in the commuter model.

These figures count the cyclists who already pass within reach of a site on the route they ride today. The model takes routes as given, so a cyclist who would make a detour to reach a service point is not counted, and the coverage is the flow that passes by without any change in behavior (Section 3.5.1).

Model	$p{=}3$	$p{=}5$	$p{=}7$	$p{=}10$
Commuter (peak)	5.9%	8.5%	11.0%	14.1%
Baseline (all-day)	5.0%	7.6%	9.8%	12.5%

Table 5.6: Share of cycling demand captured, by demand model and number of sites p. Full-population run, capture radius 150 m.

5.5 Site Structure and Cross-Model Agreement

This section reports the selected sites with the measures of Section 3.5.5. Three positions are invariant in each model: they are selected whether p is 3, 5, 7 or 10, so any solution of any size includes them.

The two models agree on 4 of their 5 sites at $p{=}5$ (within 200 m) and differ on one (Figure 5.15). The four shared sites sit on corridors used in both peak and all-day cycling,

and they include the position near Helmond, so the two demand windows agree on the regional structure of the solution as well as on the city. The single difference is local: the commuter model places a site on the Stationsweg east of Eindhoven Centraal, where the baseline model takes a position on the Nachtegaallaan 510 m to the south. In the figure each panel is one model, the grey lines are the cycle network shaded by the flow that passes along it, and the pink circles are the selected sites, numbered in the order the greedy algorithm chose them. The main map shows the part of the region where the sites fall, and the inset locates that part within the MRE.



Figure 5.15: The p=5 sites for the commuter and baseline models, numbered in selection order, over the cycle network shaded by passing flow. The inset shows the extent of each map within the MRE.

Table 5.7 lists the five commuter sites with the road they lie on, their position, their own flow and their best rival. Four sites lie in Eindhoven: on the Fellenoordpad and the Stationsweg on either side of the central station, on the Torenallee at Strijp-S, and on the Noord Brabantlaan cycleway to the west. The fifth lies on the cycleway along the Kanaaldijk Noord West in Helmond, where the route from Eindhoven meets the Zuid-Willemsvaart. The first site, on the Fellenoordpad, has the largest own flow, 2.01% of peak flow. Own flow is a per-site contribution rather than a running total, so the five values do not sum to the coverage of the set. Every selected site has a rival within 29 to 53 m and between 27 and 44 alternatives within 150 m, so the exact position is one of many near-equal points on the same corridor rather than a unique optimum.

Site	Road	Lat	Lon	Own flow	Best rival	Alternatives
1	Fellenoordpad, Eindhoven	51.4437	5.4768	2.01%	49 m	28
2	Torenallee, Eindhoven	51.4482	5.4528	1.60%	49 m	44
3	Stationsweg, Eindhoven	51.4421	5.4851	1.77%	53 m	27
4	Noord Brabantlaan, Eindhoven	51.4386	5.4318	1.33%	29 m	37
5	Kanaaldijk Noord West, Helmond	51.4736	5.6530	1.28%	50 m	38

Table 5.7: The five commuter sites (p=5), numbered in greedy selection order: road and municipality, position, own (leave-one-out) share of captured flow, distance to the best rival and number of alternatives within 150 m.

Figure 5.16 breaks the peak cycling demand down by municipality. Eindhoven dominates, carrying almost four times the flow of the next municipality, Helmond, and the five commuter sites fall in those same two municipalities. Figure 5.17 then zooms in on each of the five sites: the candidate positions around each site are coloured by the share of flow they would capture, and the chosen position (pink circle) sits on the local maximum. In every panel several neighbouring candidates reach almost the same capture, which is the

visual counterpart of the near-equivalent rivals in Table 5.7.

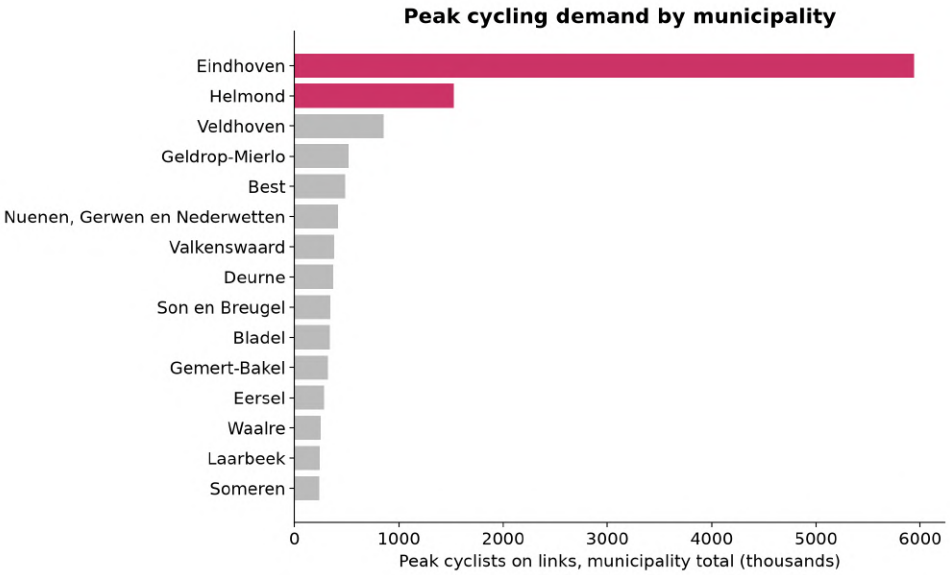

Figure 5.16: Peak cycling demand by municipality (top 15), measured as peak cyclists on the links within each municipality. Municipalities hosting one of the five commuter sites are shown in colour.

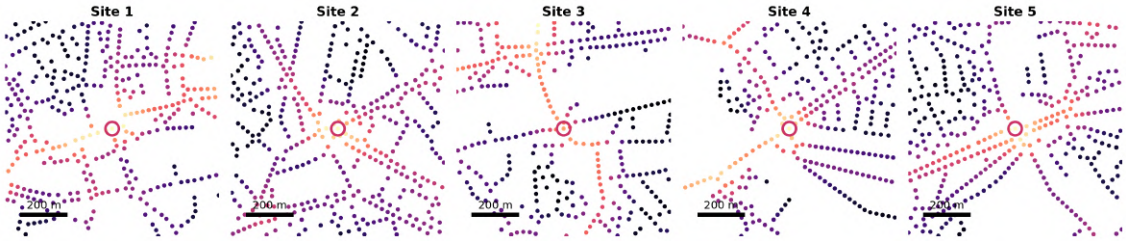

Figure 5.17: Close-up of each of the five commuter sites (p=5). Nearby candidate positions are coloured by the share of flow they would capture, and the chosen position is circled in pink.

5.6 Robustness

This section reports the four analyses of Section 3.5.6, each as the flow captured at $p = 5$ and the number of the five selected positions kept. The scenario analysis reproduces the coverage and the sites of both reported models exactly; the weight, structural and separation analyses start from the baseline model at 7.56% (Table 5.6).

5.6.1 Future Demand

Table 5.8 and Figure 5.18 give coverage at $p = 5$ for the three current scenarios and their four futures, with the sites each future shares with its current scenario. Short trips make up 86.1% of all-day and 87.1% of peak trips, so the four futures raise total demand by between 3.2% and 43.5%.

The direction of the response depends on the band and is the same in all three scenarios. Short-trip growth lowers coverage, by 0.09 to 0.19 percentage points, and long-trip growth raises it, by 0.11 to 0.34 points, so the long band moves coverage further while adding far less demand. The short-growth futures keep all five all-day sites, which gives the capture rate of each band directly: the baseline sites capture about 7.0% of short trips against 11.3% of long trips.

Sites change in eight of the twelve futures and never by more than two of five. Long-trip growth moves the Strijp-S position to Woensel-West and the Noord Brabantlaan position to the Kanaaldijk cycleway in Best in both all-day scenarios, short-trip growth moves only the Helmond commuter position and only by 50 m along the same link, and the Fellenoordpad position is selected in all fifteen runs.

	Current	Short +25%	Short +50%	Long +25%	Long +50%
S1	7.56%	7.45% (5)	7.38% (5)	7.67% (3)	7.89% (3)
S2	8.54%	8.45% (4)	8.39% (4)	8.66% (4)	8.81% (3)
S3	7.59%	7.48% (5)	7.40% (5)	7.70% (3)	7.93% (3)

Table 5.8: Flow captured at $p = 5$ for each current scenario and its four band-targeted futures. Short trips are shorter than 7.5 km and long trips are 7.5 km or longer. The number in brackets is the number of the five selected positions shared with the current scenario.

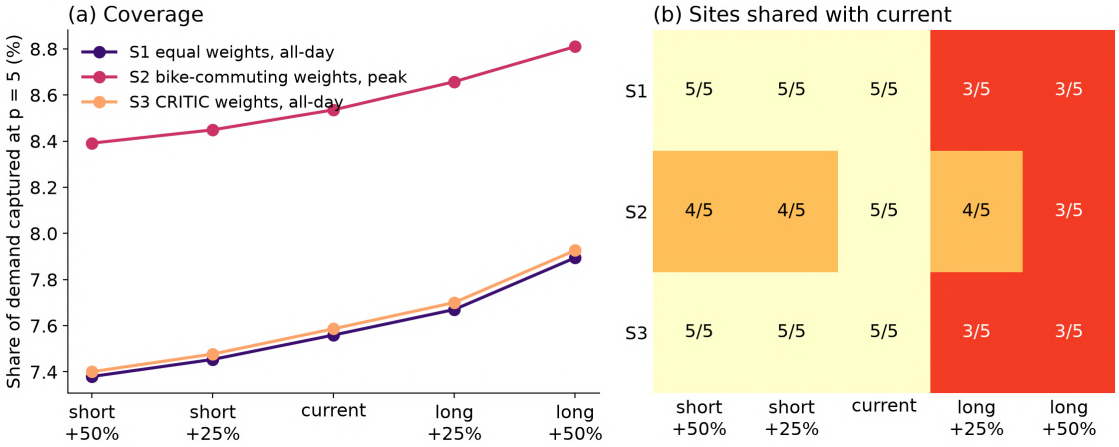

Figure 5.18: (a) Flow captured at $p = 5$ for the three current scenarios and their four band-targeted futures. (b) Number of the five selected positions each future shares with its current scenario.

5.6.2 Factor Weights

Table 5.9 gives the change in coverage when each factor weight in turn is halved and doubled from the equal profile. The largest single movement is −0.18 percentage points, from doubling land-use mix (M), and the spread across all sixteen perturbations is 0.27 points. Up-weighting M lowers coverage because its geography disagrees most with where the flow is, while up-weighting betweenness (C), business density (B), accessibility (A), main cycle network (F), transit hubs (H) or urban core (U) raises coverage by 0.01 to 0.07 points, all six pointing toward the corridors the flow already follows. Route-junction proximity (X) moves coverage by at most 0.01 points either way, being the one factor that varies at the scale of the capture kernel.

	Factor	×0.5	×2
M	land-use mix	+0.09	−0.18
X	route junction	+0.01	−0.01
U	urban core	+0.03	+0.01
H	transit hub	−0.01	+0.02
F	main cycle network	+0.01	+0.02
A	accessibility	−0.03	+0.05
B	business density	−0.03	+0.05
C	betweenness	−0.04	+0.07

Table 5.9: Change in flow captured at $p = 5$ when a single factor weight is multiplied by 0.5 or by 2 and the weight vector is renormalized, measured in percentage points against the equal-weight base of 7.56%.

Realized influence tracks the nominal weights closely: under the equal profile main cycle network reaches $w_c\sigma_c = 0.16$ and land-use mix 0.11 against a nominal 0.125. Across 100 joint weight draws coverage has a mean of 7.57% and a standard deviation of 0.08 points, and 94 draws select the same five positions within 150 m.

5.6.3 Structural Parameters

Table 5.10 reports the three structural sweeps. Coverage rises with the capture radius, from 6.98% at 75 m to 10.09% at 500 m, and the positions change with it, from three of five kept at 75 m to none at 500 m. Halving or doubling every factor buffer moves coverage by at most 0.17 points and keeps all five sites, so the reach over which the factors are computed decides nothing about where the sites go. Coarsening the candidate spacing costs 0.07 points and two sites at 50 m, and 0.35 points and all five at 100 m.

Sweep	Value	Coverage	Shared sites
Capture radius	75 m	6.98%	3/5
	150 m (base)	7.56%	5/5
	300 m	8.29%	1/5
	500 m	10.09%	0/5
Factor buffers	×0.5	7.47%	5/5
	×1 (base)	7.56%	5/5
	×2	7.73%	5/5
Candidate spacing	25 m (base)	7.56%	5/5
	50 m	7.49%	2/5
	100 m	7.21%	0/5

Table 5.10: Flow captured at $p = 5$ and the number of the five selected positions shared with the unperturbed run, for the capture radius, the factor buffer radii and the candidate spacing. Rows marked base are the values used throughout Phase 2.

5.6.4 Minimum Separation

Without a separation rule the closest pair of selected positions lies 925 m apart, against a capture kernel that reaches 150 m, so an imposed separation of 250 or 500 m changes nothing at any budget. Above that, coverage at $p = 5$ falls by 0.32 percentage points at 1 km,

by 0.69 at 2.5 km and by 1.79 at 7.5 km, and the loss grows with the budget: at 7.5 km, $p = 10$ captures 7.63% against 12.49% unconstrained.

5.7 Conclusion

This chapter established the following.

The demand object. The commuter model carries 195,513 peak trips against 235,844 candidate positions and the baseline model 626,944 all-day trips against 269,379 candidates. In both builds every trip is within reach of at least one candidate, and the single best position captures 2.06% of all-day flow on its own.

Coverage. Five service points capture 8.5% of peak flow in the commuter model and 7.6% of all-day flow in the baseline model, rising to 14.1% and 12.5% at ten. Each added site captures between 1.0 and 1.3 percentage points of peak flow across the range $p = 3$ to $p = 10$.

Site structure. Three positions in each model are selected at every budget. The two models agree on four of their five positions at $p = 5$ within 200 m and differ on one, two station-area positions 510 m apart. Every selected site has a near-equal rival within 29 to 53 m and between 27 and 44 alternative candidates within 150 m.

Factors. The maximum pairwise Spearman correlation between the eight quality factors is 0.78, and seven of the eight inter-correlate between roughly 0.24 and 0.78, with only main cycle-network proximity below 0.2 against all others. Siting on demand alone captures 9.21% of peak flow at $p = 5$, against at most 3.46% for any single factor's geography.

Robustness. No band-targeted growth future moves more than two of the five positions in any scenario. Halving or doubling a single factor weight moves coverage by at most 0.18 percentage points, and 94 of 100 Dirichlet weight draws select the same five positions within 150 m. The capture radius is the one swept parameter that moves both quantities: coverage runs from 6.98% at 75 m to 10.09% at 500 m, and no selected position survives at 500 m. The unconstrained solution already spaces its sites at least 925 m apart, so separation constraints up to 500 m leave it unchanged.

Chapter 6

Conclusions and Discussion

This chapter first answers the research question and sets out what the answer means in practice and for policy. It then evaluates the model behind that answer, reading its results against the literature, and closes with the limitations of the work and the directions they open for further research.

6.1 Conclusions

This thesis asked how a demand-informed framework can be developed to determine Cycle Service Point (CSP) locations, and where CSPs should be placed in the Metropolitan Region Eindhoven (MRE) to best support cycling commuters. It answers both halves.

The framework couples an agent-based model of cycling demand with a flow-capturing siting model. The first part turns a sparse national travel survey into a full synthetic population of cyclists, sends each of their trips over the regional cycle network, and so produces an hourly picture of where cycling runs, street by street. That picture reproduces the observed counts well enough for strategic planning, and its central finding is that cycling is concentrated on a small number of corridors. The second part chooses service-point locations against that picture: it places them where they intercept the routes cyclists actually ride, rather than near where people live, and grades each position on eight open-data criteria.

For the MRE, the answer to the where question is a small set of corridors in and around Eindhoven, with one in Helmond. The strongest positions lie on the cycle routes on either side of Eindhoven Centraal station, at Strijp-S, on the Noord Brabantlaan in western Eindhoven and on the canal-side route through Helmond where the fast cycle route from Eindhoven arrives. Peak commuting and all-day cycling pick the same four of these five and differ only over which side of the station area the fifth should serve, a choice of 510 m. Five service points on these corridors reach about 8.5% of all peak cycling flow in the region, some $16{,}700$ trips on an average Monday-to-Friday day, and each further service point adds a similar amount. Because the demand is resolved at the level of street links and every selected position has near-equal neighbors, the answer is a ranked set of corridor stretches, with the exact position left to local land availability. The answer holds under growth in cycling and under different criterion weights, and it depends above all on how far a cyclist is assumed to go out of the way for a service point.

The work makes two contributions. The academic contribution is the demand object given to the siting model. By siting against the behaviorally realized, link-by-link path of every simulated cyclist, the framework applies a flow-capturing formulation to cycling at street level on national open data. The practical contribution is a case study that delivers both a regional cycling-demand picture that is useful in its own right and a ranked set of service-point locations for the MRE.

6.2 Practical and Policy Relevance

The introduction set out a societal aim for this work: to remove practical barriers to cycling to work and so support a shift away from the car. The results translate into three implications for policy.

The first comes from the demand model. Because weekday cycling runs on a thin set of corridors, measures aimed at cyclists, whether service points, parking or route quality, reach most cyclists when they are concentrated on a limited part of the network, and the demand picture shows which part. This fits policy at both national and local level. The national Tour de Force program aims for 20% more cycled kilometers in 2027 than in 2017 and puts cycling forward as a way to relieve roads and public transport in the peak (Rijksoverheid, nd), and Eindhoven adopted its Masterplan Mobiliteit 2050 in 2024, which includes measures to complete the city's cycle network (Gemeente Eindhoven, 2024).

The second comes from the siting. The five commuter locations reach about 16,700 peak trips on an average weekday, and a CSP can only lower barriers for the cyclists it reaches, so this is the group on which the societal aim acts. The framework does not estimate how many new cyclists a service point attracts, and it does not convert coverage into fewer car trips. Evidence from the region shows the order of shift that better cycling provision on such corridors can achieve: on the F270 fast cycle route between Eindhoven and Helmond, one in five users previously made the same journey by car or public transport (Brainport Bereikbaar, 2024). The position both demand windows place in Helmond sits on that corridor, so a CSP there adds a service layer to a route the region is already investing in, at a place where a shift from the car relieves a congested road. A second such corridor enters the solution only under growth in longer trips, when a position on the canal-side route in Best is selected in both windows. That corridor lies in the area of the national MIRT study of the A2/N2 ring road west of Eindhoven, which expects traffic on the ring road to jam more often as the Brainport region grows and names fast and safe cycle routes between homes and workplaces as part of the solution (Ministerie van Infrastructuur en Waterstaat, 2026).

The third concerns the size of the investment. The Fietsservicepunt Schiphol Noord, realized by BAM, indicates an investment of about EUR 2 million per facility, and EUR 1.5 to 2.0 million is a workable planning range (BAM Infra Nederland, 2025). Five service points then represent EUR 7.5 to 10 million and ten represent EUR 15 to 20 million. Because each additional site keeps adding a similar share of peak flow, between 1.0 and 1.3 percentage points or about 2,000 to 2,500 peak trips, the framework makes the trade-off explicit for a policymaker: the order in which positions are built follows from the ranking, and the number to build follows from the budget set against the flow each further site reaches.

For practice, the recommendations follow from what the model can and cannot resolve. A municipality should build from the top of the ranking, starting with the Fellenoordpad position beside Eindhoven Centraal, which is selected under every demand window, weighting and growth future tested. No growth future moves more than two of the five positions in either model, so acting on the ranking does not have to wait for a demand forecast. The corridor whose standing does depend on the future is the one in Best: it enters the solution in both demand windows once growth lands in longer trips, as the national e-bike trend suggests it will. What a municipality does need to settle first is what a service point is meant to serve, because the walking distance assumed around each point decides the map. Each recommendation should enter local appraisal as a corridor stretch of about 150 m rather than as a single parcel. Independently of any facility decision, the demand picture itself can support network prioritization, wayfinding and safety work, because it shows where and when cycling concentrates across the region.

6.3 Discussion of the Model

The demand model is dependable on the quantities the siting depends on: how much cycling there is and which corridors carry it. Population enters the trip-generation model as an offset, so the zonal trip totals follow the number of residents, and the two retained covariates, land-use mix and cycling-lane density, add a small but significant gradient on top. This is the expected role of built-environment variables in cycling demand, where land-use mix and cycling infrastructure are consistently associated with more cycling (Cervero et al., 2009; Heinen et al., 2010). The overall level is anchored to the survey's own weights and is reproduced independently: the synthetic population matches the cyclist count and the trips per cyclist of the survey within 1%, and the 0.86 bicycle trips per inhabitant per day fall inside the Dutch metropolitan range. The purpose-specific gravity model then reproduces the observed distribution of trip lengths, in line with the entropy-maximizing formulation it follows (Wilson, 1967). Station-level studies of bicycle-sharing demand relate cycling to the same kind of built-environment variables, but they estimate on the continuous trip records of a sharing system rather than on a sample of travel diaries (Faghih-Imani and Eluru, 2016; El-Assi et al., 2017). With a diary base, the population offset carries most of the level and only two covariates survive selection, so zones differ from one another less than a larger sample would show. The corridor pattern the siting relies on comes mainly from routing the trips over the network and is less exposed to this than the zone totals.

The model is less precise about exactly which zone a trip starts in and exactly at what hour it departs. The survey provides only about 1,476 weekday bicycle trip records spread over 113 zones, so the differences between individual zones rest on a handful of diaries each, and the expansion factors vary patchily across the region rather than smoothly. In time, the simulated afternoon peak falls around 14:00 where the counts peak at 16:00. The early peak is inherited from the survey diaries, in which school trips and trips by older residents cluster in the early afternoon, while the counting stations lie mostly on routes between towns where the working-age commute dominates. For the siting this means that the demand picture can be trusted at the level of regional corridors and daily totals, and less at the level of a single street and hour.

Against the counts, the fit is acceptable for a strategic, region-scale cycling model. It misses the motor-traffic standard of Geoffrey E. Havers (GEH)<5 for 85% of counts (Highways Agency, 1996), reaching 78.1%, but 94.3% of station-hours stay below a GEH of 10. That standard was written for motorized volumes far larger than bicycle counts, and cycling has no equivalent agreed validation standard. On the measures that do transfer, the fit is good: a mean Scalable Quality Value (SQV) of 0.905 lies in the highest quality class used in agent-based model validation (Scherr et al., 2020), and the Nash–Sutcliffe efficiency of 0.518 exceeds the value of 0.5 generally taken as satisfactory for the measure (Moriasi et al., 2007).

The route choice is simpler than in the bicycle extension it builds on. Ziemke et al. (2017) let routing respond to surface, comfort, infrastructure type and gradient. Here the comfort and surface coefficients are held at zero, most links share one speed, and preference enters only through a three-level attractiveness score that favors official cycle routes. Because cyclists accept detours to ride on dedicated paths (Heinen et al., 2010), a fuller infrastructure preference would shift some flow from parallel streets onto dedicated paths outside the official route network. Flow on such paths is therefore more likely to be understated than overstated.

The one Phase 1 finding that carries structural weight for the siting is the concentration of demand: the busiest 20% of loaded links carry 80% of all link crossings. This is the condition under which a small number of well-placed points can intercept a large share of movement, and it is what makes a flow-capturing model, rather than a covering model, the appropriate choice (Hodgson, 1990).

The siting model resolves corridors, not points. The capture kernel reaches 150 m, and every selected site has a near-equal rival within 27 to 53 m, so presenting a coordinate to

the meter would claim a precision the demand does not support. A stretch of corridor is both the more honest and the more useful object for a planner. This follows directly from the gradual capture that the model adopts, in which a facility close to a route still intercepts part of its flow (Berman et al., 2003).

The design is flow-first and quality-second, and the results bear this out. Siting on demand alone captures far more flow than siting on any single quality criterion, and the best single criterion, betweenness centrality, captures little more than a third as much. This matches the logic of discretionary facilities that travellers use while already on their way, for which being on the path matters more than the character of the place (Berman et al., 1992). The quality factors refine which of several high-flow positions is chosen, and they do not decide where the flow is. Their effect is further limited by overlap: seven of the eight factors largely measure the same urban intensity, and only proximity to the main cycle network adds independent information. Because quality is a weight of at most one on captured flow, no weighted solution can capture more flow than the flow-only reference, which is what the formulation implies rather than a shortcoming of the multi-criteria layer. This separates the approach from Geographic Information System (GIS)-based siting in which a weighted sum of criteria is itself the objective, as in the genetic-algorithm model of Li and Yeh (2005). There the weights decide the location directly, whereas here they rank positions that already carry flow.

That overlap is itself partly a finding about the data. Two of the eight factors, accessibility to opportunities and business density, were first built as sums of the capacity attribute the facility register carries for every entry. The attribute is present, numeric and plausible-looking, but it does not encode one quantity across the register: the entries derived from open mapping data carry an area in square metres rather than a count of people, and a small number carry sentinel values. The consequence was that accessibility ranked rural positions next to large green areas above the city, and only a spatial check on where the opportunities were said to be exposed it. Counting each registered facility once removes the defect, and it also removes the apparent independence that accessibility had: once corrected, it correlates with urbanity as strongly as the rest. The lesson generalises beyond this model. In a multi-criteria layer, a factor that appears to add an independent dimension is worth testing against a simple spatial expectation before it is weighted, because an additive scheme has no way of noticing that one of its criteria is measuring something else.

The result also agrees with the siting logic of the CSP concept, which places facilities on high-intensity cycling corridors and at multimodal transfer nodes (Loendersloot, 2022). In the MRE the two coincide, as the strongest positions lie on the corridors on either side of Eindhoven Centraal station. Because transit proximity overlaps with the other urban-intensity factors, the model reaches these transfer nodes through the flow that passes them rather than through the transit criterion.

The two demand windows agree on four of their five positions and differ on one. That difference is local: two positions in the station area of Eindhoven $510\,\mathrm{m}$ apart, one on either side of the tracks. Peak commuting and all-day cycling therefore agree on the regional structure of the answer, including the position in Helmond, and disagree only over which approach to the station the fifth point should serve. Which of the two a region chooses depends on which cycling it wants to serve, and that is a planning decision rather than a model output. The classical flow-capturing models take their flows from origin–destination volumes routed on shortest paths (Hodgson, 1990). Siting against simulated trajectories adds a time dimension to the same objective: one simulation yields both demand windows, so the agreement between them is a result the model produces rather than an assumption, and the one place where they part can be located to the street.

Five service points capturing about 8.5% of peak cycling flow is a substantial result for point facilities spread over a regional network of $22{,}000\,\mathrm{km}$, and it follows from the concentration of demand on a few corridors. The coverage counts cyclists who already pass within reach of a site, which is the right starting point given that the mobility-hub

literature names public indifference, hubs that go largely unused, as one of the main risks of implementation (Nikitas et al., 2025). Whether the passing cyclists actually stop, and whether some of them come to cycle because the facility is there, the model does not predict. Tavassoli and Tamannaei (2020) take the other route when siting bike-and-ride hubs: they maximize the car trips shifted to the combined mode, which yields a modal-shift figure directly but requires a calibrated mode-choice model. Counting passing flow needs no such model and runs on open data alone, and it stops one step short of a modal-shift estimate.

The siting is robust to how much cycling grows and moves only modestly with what kind of cycling grows. Raising short trips by half adds 43% to all-day demand and leaves every all-day position in place, while raising long trips by half adds only 7% and changes two of the five. Longer trips pass more candidate positions and run along the corridors between towns, so they are captured at a clearly higher rate, about 11% against 7% for short trips, and growth in long trips raises coverage while growth in short trips lowers it slightly. No future tested changes more than two of the five sites in either model. What moves is consistent across both: whenever long trips grow, the canal-side corridor in Best is gained, which is the direction the national projections lean, with e-bike distance expected to grow while conventional cycling declines (Huang et al., 2024). Short-trip growth moves nothing in the all-day models, and in the commuter model it moves the Helmond point by 50 m along the same link, which leaves the place unchanged and only the exact point open. The Fellenoordpad position is selected in every run.

The criterion weights matter less than their prominence in the method suggests. Halving or doubling any single weight moves the captured flow by less than a fifth of a percentage point, 94 of 100 random weightings around equal weights select the same five positions, and weights derived objectively from the contrast and redundancy of the factor scores (Diakoulaki et al., 1995) arrive at the same positions but one, and that one is a neighbouring candidate 49 m away. The single-author weighting, the most subjective part of the method, therefore has the least leverage over the answer. A mobility expert at the Province of Noord-Brabant reviewed the weights and judged them sound, which gives the practitioner reading an outside check without the pairwise elicitation of a full panel.

The capture radius is the one assumption that does carry the result. Widening it from 150 to 500 m raises coverage and replaces every selected position. This is not instability, because the radius is a statement of what a service point is assumed to serve: a short radius asks where to intercept cyclists who pass close to the door, and a wide radius asks where to serve a catchment. It is a service-design decision to be taken before the optimization. Finally, the objective spaces its sites naturally. The selected positions lie at least 925 m apart, so the model needs no separation rule, and imposing one above a kilometer only removes captured flow.

The methodological contribution lies in the demand object rather than in the optimizer, which is a standard greedy coverage solver with a known performance guarantee (Nemhauser et al., 1978). Existing siting work tends to cover residents, zonal population or static suitability surfaces. Here the sites are chosen against the realized path of each simulated agent, built with an established agent-based simulation platform (Horni et al., 2016), so the model knows whether two links carry the same cyclist and counts each trip once. A smaller supporting element is that the accessibility factor's decay is derived from the simulation's own travel utility, so it re-derives wherever the demand model is re-run.

The framework transfers in its mechanics but not in its weights. Every input is a nationally available open-data source and every step uses open-source software, so applying it to another Dutch urban region means substituting the regional layers and re-running both phases. The larger metropolitan regions around Rotterdam, The Hague and Amsterdam are the natural next cases, since they have the survey coverage and the cycling volumes the method assumes. The criterion weights encode one practitioner's judgement and would need re-eliciting for a region with different priorities. Two conditions bound the transfer: the demand model needs enough survey respondents per region to estimate zonal rates,

and outside the Netherlands the open-data inputs and the strong cycling base cannot be assumed, so transfer abroad is a matter of re-sourcing inputs rather than re-running a script.

6.4 Limitations and Future Work

The binding limitation on the demand side is spatial resolution. The four-digit postcode zones used for trip generation are coarser than a typical cycling trip of a few kilometers, so fine-grained variation in where demand originates cannot be resolved regardless of how large the simulated population is, and the siting result inherits that ceiling. This sits on top of a thin estimation base of about 1,476 weekday trip records across 113 observed zones, which is what widens the zone-level scatter and leaves the expansion factors patchy.

The temporal fit is partial. The model reproduces the morning peak but places the afternoon peak around 14:00 rather than 16:00, because departure times are copied from the survey diaries and trip distance is assigned without regard to departure hour. The peak-hour demand the commuter model intercepts is therefore right in total but not perfectly timed in the afternoon.

A residual 5.1% of planned trips are not executed because the network cannot clear them, an effect of the storage-capacity floor on an unusually short-linked cycling network. The loss is small and falls on the most congested links, so it slightly understates flow where flow is heaviest.

Validation is partial and one-sided. The demand surface can be checked against permanent counters, but the siting result has no external register of realized CSP locations to validate against, because the facility concept is new and almost none have been built. The siting model is therefore internally consistent but externally untested, and its output should be read as strategic guidance rather than as a validated prediction.

The criterion weights are a single-author elicitation. They rest on a practitioner reading of what makes a good position and were reviewed by one expert at the province, but were not drawn from a panel, and the alternative weight setups the expert suggested were not run. Because seven of the eight factors also overlap strongly, the effective number of distinct criteria is smaller than eight, and the weighting has less room to change the outcome than its apparent richness suggests. Betweenness centrality is computed as a node-based approximation within a 7.5 km cutoff rather than as true edge betweenness, which is computationally prohibitive on a network of this size.

Two of those factors also lost a dimension when the capacity attribute of the facility register turned out to be unusable. Business density and accessibility now count facilities rather than weighting them by size, so a position near ten small employers scores the same as one near ten large ones. The register carries no other size measure that means the same thing for every entry, and building floor area from the address and buildings register would be the honest replacement. Until it is added, both factors describe how many opportunities are nearby and not how much demand those opportunities represent.

Coverage measures the cycling flow that passes within reach of a service point, not the number of cyclists who will use it or switch from the car because of it. The model also takes routes as given and does not credit cyclists who would make a detour to reach a facility. The uptake and modal-shift effect at the selected positions therefore remain to be observed once facilities are in operation.

The capture radius is a stated service-design assumption rather than an estimated quantity, and it is the one input the siting result turns on: at 500 m none of the five positions survive. The 150 m used throughout is a defensible reading of the walking distance a passing cyclist would accept for a quick service stop, but it was not observed, and a different service concept would produce a different map.

The future-demand scenarios are bounded in two ways. They reweight the trips the demand model produced, so they can change how much a corridor is worth but cannot put

cycling on a corridor that carries none today. And the growth levels per trip-length band are exploratory magnitudes rather than forecasts. The finding is therefore that the siting is robust to growth in the volume of demand on the existing network, which is a narrower claim than robustness to the region's cycling future in general.

The analysis is bounded to weekday utility with a strong focus on commuter cycling. This matches the CSP concept, which serves commuters, but it means the results speak only to that segment. Buildability enters as a check after siting rather than as a filter inside candidate generation. This is by design, because the aim was to identify strong locations at corridor level, but a selected corridor still has to pass a site-level feasibility test before anything is built.

Several extensions follow from these limitations. Running the alternative weight setups the provincial expert suggested, or eliciting the weights from a small practitioner panel, would put the multi-criteria layer on a collective footing. Lifting the demand resolution below the postcode zone, through a finer zoning or a disaggregate residential allocation, would sharpen where demand originates. Once a register of realized CSPs exists, validating the siting model against it would let its second half be tested externally.

Two extensions follow from the robustness results. The first is a re-simulated future: putting a grown or redistributed population back through the demand model, for example around a planned housing development, along a new cycle highway or with the projected shift toward longer e-bike trips, so that growth can reach corridors the region does not yet ride. The second is to settle the capture radius empirically with a small revealed or stated preference study of how far a cyclist will divert for a quick service stop, which would turn the most consequential assumption of the siting model from a choice into an estimate.

Further refinements concern the inputs and the factors. Reconciling the network with building-footprint and open-space data would let buildability move into candidate generation. Joining building floor area from the address and buildings register to the facility points would give business density and accessibility back the size dimension the unusable capacity attribute cost them. Harmonizing the factor kernels, two of which still use a hard cutoff at their reach, and computing betweenness on edges rather than nodes would make the quality layer more consistent. Collecting cycling data beyond the travel survey, through counts, GPS traces or app-based flows, would reduce the dependence on a small diary sample. Together these steps would move the framework from strategic guidance toward validated, buildable recommendations.

Bibliography

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723. https://doi.org/10.1109/TAC.1974.1100705. 13, 30

BAM Infra Nederland (2025). Start realisatie fietsservicepunt bij busstation Schiphol Noord. [News release]. Realisation cost range from BAM project information. Retrieved September 14, 2026, from https://www.baminfra.nl/nieuws/2025/12/start-realisatie-fietsservicepunt-bij-busstation-schiphol-noord. 2, 9, 90

Barthelemy, M. (2016). A global take on congestion in urban areas. *Environment and Planning B: Planning and Design*, 43:800–804. https://doi.org/10.1177/0265813516649955. 1, 17

Bergström, A. and Magnusson, R. (2003). Potential of transferring car trips to bicycle during winter. *Transportation Research Part A: Policy and Practice*, 37(8):649–666. https://doi.org/10.1016/S0965-8564(03)00012-0. 7

Berman, O., Krass, D., and Drezner, Z. (2003). The gradual covering decay location problem on a network. *European Journal of Operational Research*, 151(3):474–480. https://doi.org/10.1016/S0377-2217(02)00604-5. 16, 92

Berman, O., Larson, R. C., and Fouska, N. (1992). Optimal location of discretionary service facilities. *Transportation Science*, 26(3):201–211. https://doi.org/10.1287/trsc.26.3.201. 10, 16, 18, 92

Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139. https://doi.org/10.1016/j.compenvurbsys.2017.05.004. 17, 33

Brainport Bereikbaar (2024). Tijdwinst op de fiets tussen Helmond en Eindhoven. [News release]. Retrieved September 14, 2026, from https://www.brainportbereikbaar.nl/en/inspiratie/tijdwinst-op-de-fiets-tussen-helmond-en-eindhoven/. 90

Cameron, A. C. and Trivedi, P. K. (2013). *Regression Analysis of Count Data*. Cambridge University Press, Cambridge, 2nd edition. https://doi.org/10.1017/CBO9781139013567. 13, 30, 54

Centraal Bureau voor de Statistiek (2017). Bestand bodemgebruik 2015. [Data set]. https://www.pdok.nl. 23

Centraal Bureau voor de Statistiek (2022a). Kerncijfers per postcodegebied (PC4), 2022. [Data set]. https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/gegevens-per-postcode. 23

Centraal Bureau voor de Statistiek (2022b). Wijk- en buurtkaart 2022. [Data set]. https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/wijk-en-buurtkaart-2022. 23

Centraal Bureau voor de Statistiek (2023). Onderweg in Nederland (ODiN) 2022. [Data set]. Commissioned by Rijkswaterstaat. https://doi.org/10.17026/SS/BXIK2X. 23, 26

Cervero, R. and Kockelman, K. (1997). Travel demand and the 3Ds: Density, diversity, and design. *Transportation Research Part D: Transport and Environment*, 2(3):199–219. https://doi.org/10.1016/S1361-9209(97)00009-6. 43

Cervero, R., Sarmiento, O. L., Jacoby, E., Gomez, L. F., and Neiman, A. (2009). Influences of built environments on walking and cycling: Lessons from Bogotá. *International Journal of Sustainable Transportation*, 3(4):203–226. https://doi.org/10.1080/15568310802178314. 6, 30, 91

Chang, Y. S., Lee, Y. J., and Choi, S. S. B. (2017). Is there more traffic congestion in larger cities? Scaling analysis of the 101 largest U.S. urban centers. *Transport Policy*, 59:54–63. https://doi.org/10.1016/j.tranpol.2017.07.002. 1

Charypar, D. and Nagel, K. (2005). Generating complete all-day activity plans with genetic algorithms. *Transportation*, 32(4):369–397. https://doi.org/10.1007/s11116-004-8287-y. 14, 24, 35

Church, R. and ReVelle, C. (1974). The maximal covering location problem. *Papers of the Regional Science Association*, 32:101–118. https://doi.org/10.1007/BF01942293. 15

CyclOSM contributors (2024). CyclOSM: A bicycle-oriented map rendering of OpenStreetMap. https://www.cyclosm.org. Bicycle access tag vocabulary (`can_bicycle`). 27

de Dios Ortúzar, J. and Willumsen, L. G. (2011). *Modelling Transport*. Wiley, Chichester, 4th edition. https://doi.org/10.1002/9781119993308. 13, 24, 37

de Kruijf, J., Ettema, D., Kamphuis, C. B. M., and Dijst, M. (2018). Evaluation of an incentive program to stimulate the shift from car commuting to E-cycling in the Netherlands. *Journal of Transport & Health*, 10:74–83. https://doi.org/10.1016/j.jth.2018.06.003. 1, 2, 7, 8, 9

Deb, K. (2001). *Multi-Objective Optimization Using Evolutionary Algorithms*. Wiley, Chichester. 16

Diakoulaki, D., Mavrotas, G., and Papayannakis, L. (1995). Determining objective weights in multiple criteria problems: The CRITIC method. *Computers & Operations Research*, 22(7):763–770. https://doi.org/10.1016/0305-0548(94)00059-H. 17, 47, 93

Drezner, Z. and Hamacher, H. W., editors (2002). *Facility Location: Applications and Theory*. Springer, Berlin. https://doi.org/10.1007/978-3-642-56082-8. 15, 16

El-Assi, W., Mahmoud, M. S., and Habib, K. N. (2017). Effects of built environment and weather on bike sharing demand: A station level analysis of commercial bike sharing in Toronto. *Transportation*, 44(3):589–613. https://doi.org/10.1007/s11116-015-9669-z. 7, 13, 30, 91

Faghih-Imani, A. and Eluru, N. (2016). Incorporating the impact of spatio-temporal interactions on bicycle sharing system demand: A case study of New York Citibike system. *Journal of Transport Geography*, 54:218–227. https://doi.org/10.1016/j.jtrangeo.2016.06.008. 13, 91

Fishman, E., Washington, S., and Haworth, N. (2012). Barriers and facilitators to public bicycle scheme use: A qualitative approach. *Transportation Research Part F: Traffic Psychology and Behaviour*, 15(6):686–698. https://doi.org/10.1016/j.trf.2012.08.002. 2

Friedrich, M., Pestel, E., Schiller, C., and Simon, R. (2019). Scalable GEH: A quality measure for comparing observed and modeled single values in a travel demand model validation. *Transportation Research Record*, 2673(4):722–732. https://doi.org/10.1177/0361198119838849. 15, 38, 107

Gemeente Eindhoven (2024). Masterplan mobiliteit 2050. [Policy document]. Adopted by the municipal council on March 12, 2024. Retrieved September 14, 2026, from https://www.openeindhoven.nl/project/masterplan-mobiliteit. 90

Gemeente Eindhoven (n.d.). Hubs voor deelvervoer. [Web page]. Retrieved September 17, 2026, from https://www.eindhoven.nl/stad-en-wonen/verkeer/deelvervoer/hubs. 2

Gen, M. and Cheng, R. (1997). *Genetic Algorithms and Engineering Design*. Wiley, New York. https://doi.org/10.1002/9780470172254. 15, 16

Gupta, H. V., Kling, H., Yilmaz, K. K., and Martinez, G. F. (2009). Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. *Journal of Hydrology*, 377(1–2):80–91. https://doi.org/10.1016/j.jhydrol.2009.08.003. 107

Heinen, E., Maat, K., and van Wee, B. (2013). The effect of work-related factors on the bicycle commute mode choice in the Netherlands. *Transportation*, 40(1):23–43. https://doi.org/10.1007/s11116-012-9399-4. 2, 8, 9, 17

Heinen, E., van Wee, B., and Maat, K. (2010). Commuting by bicycle: An overview of the literature. *Transport Reviews*, 30(1):59–96. https://doi.org/10.1080/01441640903187001. 1, 2, 6, 7, 8, 11, 14, 91

Highways Agency (1996). Design manual for roads and bridges, volume 12, section 2, part 1: Traffic appraisal in urban areas. Technical report, The Stationery Office, London. 15, 37, 38, 68, 91

Hodgson, M. J. (1990). A flow-capturing location-allocation model. *Geographical Analysis*, 22(3):270–279. https://doi.org/10.1111/j.1538-4632.1990.tb00210.x. 10, 15, 16, 18, 91, 92

Horni, A., Nagel, K., and Axhausen, K. W., editors (2016). *The Multi-Agent Transport Simulation MATSim*. Ubiquity Press, London. https://doi.org/10.5334/baw. 13, 14, 15, 24, 35, 93

Huang, B., Wüst, H., and de Haas, M. (2024). Assessing the E-bike trends and impact on sustainable mobility: A national-level study in the Netherlands. *Journal of Cycling and Micromobility Research*, 2:100027. https://doi.org/10.1016/j.jcmr.2024.100027. 2, 48, 93

Hull, A. and O'Holleran, C. (2014). Bicycle infrastructure: Can good design encourage cycling? *Urban, Planning and Transport Research*, 2(1):369–406. https://doi.org/10.1080/21650020.2014.955210. 1, 8, 11

Kadaster (2024). Basisregistratie adressen en gebouwen (BAG). [Data set]. Retrieved August 1, 2026, from https://www.pdok.nl. 23

Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., VanBriesen, J., and Glance, N. (2007). Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '07)*, pages 420–429. https://doi.org/10.1145/1281192.1281239. 16, 46

Li, X. and Yeh, A. G.-O. (2005). Integration of genetic algorithms and GIS for optimal location search. *International Journal of Geographical Information Science*, 19(5):581–601. https://doi.org/10.1080/13658810500032388. 16, 17, 18, 92

Loendersloot, P. (2022). Fietsservicepunten metropoolregio Amsterdam: Het steunpunt voor de fietsforens. 2, 7, 9, 10, 43, 92

Metropoolregio Eindhoven (n.d.). Thema mobiliteit. [Web page]. Retrieved September 17, 2026, from https://metropoolregioeindhoven.nl/thema-s/mobiliteit. 2

Ministerie van Infrastructuur en Waterstaat (2026). MIRT-verkenning A2/N2 randweg Eindhoven. [Project page]. Retrieved September 14, 2026, from https://www.mirtoverzicht.nl/lopende-mirt-projecten/mirt-verkenning-a2-n2--randweg-eindhoven. 90

Minoux, M. (1978). Accelerated greedy algorithms for maximizing submodular set functions. In *Optimization Techniques*, volume 7 of *Lecture Notes in Control and Information Sciences*, pages 234–243. Springer, Berlin, Heidelberg. https://doi.org/10.1007/BFb0006528. 16, 46

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2):17–23. https://doi.org/10.2307/2332142. 30

Moriasi, D. N., Arnold, J. G., Van Liew, M. W., Bingner, R. L., Harmel, R. D., and Veith, T. L. (2007). Model evaluation guidelines for systematic quantification of accuracy in watershed simulations. *Transactions of the ASABE*, 50(3):885–900. https://doi.org/10.13031/2013.23153. 68, 91

Nash, J. E. and Sutcliffe, J. V. (1970). River flow forecasting through conceptual models part I — a discussion of principles. *Journal of Hydrology*, 10(3):282–290. https://doi.org/10.1016/0022-1694(70)90255-6. 37

Nationaal Dataportaal Wegverkeer (2024). Fietstellingen. [Data set]. Retrieved August 1, 2026, from https://www.ndw.nu. 23

Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294. https://doi.org/10.1007/BF01588971. 12, 16, 18, 39, 46, 93

Nikitas, A., Alyavina, E., Sadik, C., and Michalakopoulou, K. (2025). Decoding mobility hubs: Opportunities and risks underpinning their introduction for the contexts of transport and the wider society. *Journal of Transport Geography*, 127:104296. https://doi.org/10.1016/j.jtrangeo.2025.104296. 2, 10, 11, 18, 93

OpenStreetMap contributors (2024). Openstreetmap. [Data set]. Retrieved August 1, 2026, from https://www.openstreetmap.org. 23

Parkin, J., Wardman, M., and Page, M. (2008). Estimation of the determinants of bicycle mode share for the journeys to work using census data. *Transportation*, 35(1):93–109. https://doi.org/10.1007/s11116-007-9137-5. 2, 7

Porta, S., Latora, V., Wang, F., Strano, E., Cardillo, A., Scellato, S., Iacoviello, V., and Messora, R. (2009). Street centrality and densities of retail and services in Bologna, Italy. *Environment and Planning B: Planning and Design*, 36(3):450–465. https://doi.org/10.1068/b34098. 17, 42

Rijksoverheid (n.d.). Kabinet: meer mensen op de fiets. [Web page]. Retrieved September 14, 2026, from https://www.rijksoverheid.nl/onderwerpen/fiets/fietsbeleid. 90

Scheepers, E., Wendel-Vos, W., van Kempen, E., Int Panis, L., Maas, J., Stipdonk, H., Moerman, M., den Hertog, F., Staatsen, B., van Wesemael, P., and Schuit, J. (2013). Personal and environmental characteristics associated with choice of active transport modes versus car use for different trip purposes of trips up to 7.5 kilometers in The Netherlands. *PLoS ONE*, 8(9):e73105. https://doi.org/10.1371/journal.pone.0073105. 48

Scherr, W., Manser, P., Joshi, C., Frischknecht, N., and Métrailler, D. (2020). Towards agent-based travel demand simulation across all mobility choices: The role of balancing preferences and constraints. *European Journal of Transport and Infrastructure Research*, 20(4):152–172. https://doi.org/10.18757/ejtir.2020.20.4.4463. 15, 38, 68, 91

Schuhmacher, L., Wilkes, G., Kagerbauer, M., and Vortisch, P. (2025). Cycling speeds in urban traffic. *Findings*. https://doi.org/10.32866/001c.141204. 28, 101

Sevtsuk, A. (2014). Location and agglomeration: The distribution of retail and food businesses in dense urban environments. *Journal of Planning Education and Research*, 34(4):374–393. https://doi.org/10.1177/0739456X14550401. 17, 42

Sheller, M. and Urry, J. (2000). The city and the car. *International Journal of Urban and Regional Research*, 24(4):737–757. https://doi.org/10.1111/1468-2427.00276. 1, 11

Stichting LISA (2023). LISA vestigingenregister. [Data set]. https://www.lisa.nl. 23

Tavassoli, K. and Tamannaei, M. (2020). Hub network design for integrated bike-and-ride services: A competitive approach to reducing automobile dependence. *Journal of Cleaner Production*, 248:119247. https://doi.org/10.1016/j.jclepro.2019.119247. 2, 12, 18, 93

United Nations (2018). World urbanization prospects: The 2018 revision. Report ST/ESA/SER.A/420, United Nations Department of Economic and Social Affairs (UN DESA). https://doi.org/10.18356/b9e995fe-en. 1

Wilson, A. G. (1967). A statistical theory of spatial distribution models. *Transportation Research*, 1(3):253–269. https://doi.org/10.1016/0041-1647(67)90035-4. 24, 34, 91

Ziemke, D., Metzler, S., and Nagel, K. (2017). Modeling bicycle traffic in an agent-based transport simulation. In *Procedia Computer Science*, volume 109, pages 923–928. https://doi.org/10.1016/j.procs.2017.05.424. 13, 14, 15, 24, 28, 35, 91

Appendix A

Supporting Tables

This appendix collects the larger reference and parameter tables underlying the Phase 1 methodology of Chapter 3. They are reproduced here rather than in the main text because each is consulted for a specific value rather than read as part of the argument.

A.1 Observed Cycling Speeds by Infrastructure Type

Table A.1 gives the empirical basis for the free-flow speed used in the network build (Section 3.4.3).

Infrastructure type	Median speed (km/h)
Pedestrian zone	15.0
Sidewalk (walking bicycles only)	16.6
Traffic-calmed street	17.0
Bicycle path (off-street)	18.5
Shared roadway ($\leq$ 30 km/h)	19.3
Sidewalk (bicycles allowed)	20.3
Bicycle lane (on-street, marked)	20.9
Bicycle street	21.0
Shared-use path	21.0
Advisory bicycle lane	21.1
Farm road	21.6
Shared roadway ($\geq$ 50 km/h)	21.7

Table A.1: Weighted median cycling speeds by infrastructure type, from a regression over 91,083 segments of GPS-tracked cycling (Schuhmacher et al., 2025).

A.2 Trip Generation Model Diagnostics

Tables A.2 and A.3 support the trip generation results of Section 4.1. A variance inflation factor above 10 is conventionally taken to indicate problematic multicollinearity, and the fold-to-fold spread of R^2 in Table A.3 is why cross-validated R^2 is not used as a selection criterion.

Candidate predictor	VIF
pop_density	10.22
cycling_lane_density	5.46
employment_density	5.33
car_ownership_rate	3.01
land_use_mix	2.52
share_18_65	2.37
median_income	1.40

Table A.2: Variance inflation factors for the seven candidate predictors, before stepwise selection.

Fold	RMSE	MAE	R^2	Test zones
1	2,322	1,522	−0.134	19
2	2,545	2,037	0.490	19
3	3,195	1,908	−0.209	19
4	2,389	1,810	0.568	19
5	2,566	1,739	−0.498	19
6	3,183	1,948	−0.796	18
Mean	2,700	1,827	−0.097	113
SD	390	183	0.539	–

Table A.3: Six-fold cross-validation of the NB2 trip generation model.

A.3 Population Synthesis Configuration

Table A.4 lists the configuration of the population synthesis step (Section 3.4.5) for the reported run.

Parameter	Value	Role
VARIANT	monfri	Modeled day: Monday–Friday
POPULATION_SCALE_FACTOR	1.0	Full population
DEMAND_ANCHOR	trips	Anchor of Eq. (3.10)
OWN_ZONE_WEIGHT	5	λ in Eq. (3.6)
DAY_SHIFT_SIGMA_MIN	30 min	σ of the whole-day time shift
WORK_ATTRACTION	raw	Work attraction basis
TRIP_GEO_FILTER	both_ends	Both trip ends inside the region

Table A.4: Population synthesis configuration for the reported run.

A.4 Facility Activity-Type Mapping

Table A.5 gives the correspondence between the source data tags and the activity vocabulary used by the facility layer (Section 3.4.6). The final column is the fallback footprint applied when a facility has no recorded area, which arises for point-only OpenStreetMap features.

Activity	OSM tags	BAG gebruiksdoel	Fallback (m^2)
work	`office=*`; `building=office/commercial/industrial`; `landuse=commercial/industrial`	kantoorfunctie, industriefunctie	120
education	`amenity=school/university/college/kindergarten`; `building=school`	onderwijsfunctie	600
shop	`shop=*` (retail subtypes); `building=retail`; `landuse=retail`	winkelfunctie	80
leisure	`leisure=*`; `amenity=restaurant/cafe/bar/cinema/theatre/library`	sportfunctie, logiesfunctie, bijeenkomstfunctie	150
personal	`amenity=bank/post_office/townhall/social_facility`; personal-care shop subtypes	–	60
other	any remaining `amenity=*`	unmapped use classes	80
home	–	woonfunctie	100

Table A.5: Mapping of OSM tags and BAG use classes onto the activity vocabulary shared with the ODiN trip purposes, with the fallback footprint used when no real area is available.

A.5 Distance-Decay Calibration

This appendix expands on the calibration of β_p and the construction of the candidate destination set (Section 3.4.6).

Starting from $\beta_p^{(0)}$ (Equation (3.13)), destinations are assigned with the current β_p, the resulting mean assigned distance is compared against $\bar{d}_p^{\text{obs}}$, and β_p is adjusted accordingly. A purpose whose assigned mean lies within 5% of its target is left unchanged. Iteration stops once no β_p changes by more than 5% between runs, or after the cap of five assignment runs. The reported run stopped after three runs, when no β_p changed. Purposes that fail to converge are logged explicitly rather than silently accepted; a β_p that moves monotonically across iterations without settling indicates a purpose that conflates genuinely different trip types, whose distance distribution is not adequately described by a single exponential.

Onderweg in Nederland (ODiN) reports network distance (`AfstV`), whereas Equation (3.11) operates on straight-line distance d_{of}. Calibrating β_p against network distance and then applying it to straight-line distance would place every assigned destination one full detour too far from its origin. Every observed distance used as a calibration target, d^{ODiN}, is therefore divided by the network detour factor $\kappa > 1$ (Equation (3.14)) before it is used as d^{target}, so that the calibration target and the model quantity are the same kind of distance. Assigned straight-line distances consequently sit systematically below the raw ODiN means by a factor $1/\kappa$; this is the correction working as intended and must not be read as under-assignment. The distances actually ridden by the agents, measured over the network after assignment (Section 3.4.7), are the quantity comparable with `AfstV`.

The candidate set $\mathscr{C}_p(o)$ is drawn from a radius r_p derived from the calibrated decay itself (Equation (3.15)), set to capture $q = 0.99$ of the total decay mass so that facilities beyond it would contribute negligibly to Equation (3.12). Where the resulting pool exceeds a cap of 3,000 facilities, a random subsample of the pool is retained. Keeping the nearest facilities instead would re-truncate the candidate set after the radius had deliberately widened it, and so re-impose exactly the short-distance bias that Equation (3.15) exists to avoid.

A.6 MATSim Simulation Control Parameters

Table A.6 lists the full configuration of the reported Multi-Agent Transport Simulation (MATSim) run (Section 3.4.7).

Parameter	Value	Role
Simulation control		
`lastIteration`	100	Length of the co-evolutionary loop
`learningRate`	0.1	Weight of the newest score in the plan memory
`brainExpBeta`	1.0	Sensitivity of plan selection to score differences
`ReRoute`	0.20	Share of agents re-routing per iteration
`TimeAllocationMutator`	0.05	Share of agents mutating departure times
`disableInnovationAfter`	80	Applied to both innovative strategies
`routingRandomness`	3.0	Spread of the randomised least-cost path search
`randomSeed`	24	Fixed and reported; run-to-run spread is non-trivial
Network and vehicle		
`inputNetworkFile`	flat-speed	Standard ridable links at one cruising speed (Sec. 3.4.3)
`maxBicycleSpeedForRouting`	6.03 m/s	Routing speed cap; does not bind at 21.1 km/h
`bikeTravelUtil`	-7.0	Marginal utility of bicycle travel time
`attractivenessUtil`	-0.0008	Coefficient c_{marg} of Eq. (3.3)
`effectiveCellSize`	1.0 m	Space per queued bicycle (storage), Sec. 4.4.1
Scale		
`qsim endTime`	30:00:00	Allows post-midnight arrivals to complete
`flowCapacityFactor`	1.0	Full population simulated
`storageCapacityFactor`	1.0	Full population simulated

Table A.6: MATSim simulation control parameters for the reported run.

A.7 Additional Validation Metrics

This appendix defines the metrics reported alongside Geoffrey E. Havers (GEH) and Nash–Sutcliffe efficiency (NSE) (Section 3.4.8) for cross-checking.

The aggregate ratio $\beta = \bar{M}/\bar{C}$, the ratio of mean simulated volume $\bar{M}$ to mean observed volume $\bar{C}$ over all stations and hours, measures whether the model produces the right amount of cycling overall, but says nothing about whether that cycling is in the right place.

Spearman's ρ measures whether the simulated and observed series rank station-hours in the same order, and is the more robust correlation measure when a handful of stations carry disproportionate volume. Like r, it is computed on the hourly series and, separately, on daily station totals.

The variance ratio $\alpha = \sigma_M/\sigma_C$, the ratio of the simulated to the observed standard deviation σ across station-hours, states whether the model spreads demand too widely or too narrowly. It is reported explicitly because several of the other metrics can be improved simply by shrinking simulated variance towards the mean, which is not a modeling improvement, and α is what exposes that.

The Scalable Quality Value (SQV) (Friedrich et al., 2019) complements GEH by normalizing the deviation and mapping it into $(0, 1]$,

$$\mathrm{SQV} = \frac{1}{1 + \sqrt{\dfrac{(M-C)^2}{f\,C}}}, \tag{A.1}$$

where f is a scaling factor matched to the magnitude of the counts (here $f = 1000$ for the hourly counts), following Friedrich et al. (2019), and $\mathrm{SQV} = 1$ denotes a perfect match. Unlike GEH, SQV is bounded and directly averageable across stations, which makes it the more informative summary statistic when the station set is heterogeneous. It is reported both unweighted and weighted by observed volume, so that agreement at the busy stations that dominate the demand surface can be distinguished from agreement at the quiet ones.

The Kling–Gupta efficiency (KGE) (Gupta et al., 2009) was introduced because NSE conflates correlation, variance and bias into one number. It combines them explicitly,

$$\mathrm{KGE} = 1 - \sqrt{(r-1)^2 + (\alpha-1)^2 + (\beta-1)^2}, \tag{A.2}$$

with r the Pearson correlation, α the variance ratio and β the aggregate ratio defined above. Its value is diagnostic as well as summary: a poor KGE can be attributed to whichever of the three terms is furthest from unity.

Table A.7 lists the per-station fit for all 47 Nationaal Dataportaal Wegverkeer (NDW) counting stations, ordered by GEH. Its observed and simulated columns are average-weekday daily totals, whereas its GEH and SQV columns are computed per station-hour and then averaged over the day, so the latter two cannot be reproduced directly from the former two.

Station	Obs	Sim	GEH	SQV
PNB06_A12	1075	3635	9.51	0.71
PNB06_C62	816	2764	8.35	0.74
PNB06_B18	459	1665	6.82	0.75
PNB06_C60	1795	3358	6.29	0.81
PNB06_B02	1289	300	6.25	0.87
PNB09_D01	720	30	6.02	0.89
PNB09_D33	990	275	5.32	0.89
PNB06_C42	2604	1590	5.09	0.88
PNB09_D31	1074	416	4.55	0.90
PNB09_D20	1568	814	4.46	0.90
PNB06_A11	2112	2231	4.12	0.89
PNB09_D23	2827	1848	3.90	0.91
PNB09_D14	2898	2505	3.83	0.90
PNB09_D05	1217	1730	3.56	0.90
PNB06_A14	549	1061	3.43	0.88
PNB09_D13	304	47	3.41	0.92
PNB09_D15	4136	4540	3.25	0.91
PNB09_D108	833	1422	3.24	0.90
PNB06_C45	1137	1626	3.13	0.90
PNB06_C63	220	529	3.02	0.89
PNB06_C61	742	943	2.97	0.90
PNB09_D07	442	161	2.95	0.93
PNB09_D02	485	869	2.90	0.89
PNB09_D18	1399	1829	2.87	0.91
PNB09_D17	1117	639	2.86	0.93
PNB09_D34	1424	1503	2.85	0.92
PNB09_D11	1946	2432	2.79	0.92
PNB09_D16	2718	2174	2.70	0.93
PNB06_C46	989	1312	2.50	0.92
PNB09_D22	1422	1169	2.45	0.93
PNB09_D24	855	581	2.42	0.94
PNB06_B03	553	578	2.40	0.93
PNB06_B05	1008	1082	2.38	0.93
PNB09_D12	169	29	2.36	0.94
PNB06_B04	1013	891	2.34	0.93
PNB09_D04	2051	2102	2.20	0.94
PNB06_B01	528	321	2.17	0.95
PNB09_D100	321	156	2.06	0.95
PNB09_D03	1071	916	2.03	0.94
PNB09_D09	1724	1793	1.96	0.94
PNB06_C43	855	1051	1.84	0.94
PNB09_D102	225	133	1.82	0.95
PNB06_C64	162	230	1.71	0.94
PNB06_C44	1652	1715	1.68	0.95
PNB09_D08	639	646	1.63	0.95
PNB09_D107	292	450	1.59	0.94
PNB06_C41	236	229	0.97	0.97

Table A.7: Observed and simulated average-weekday volume, GEH and SQV for each of the 47 NDW counting stations, ordered by GEH (worst first).

A.8 Per-Factor Suitability Surfaces

Full-size versions of the eight quality-factor surfaces of Section 5.3, each normalized to $[0, 1]$ and drawn at candidate resolution over the all-day build.

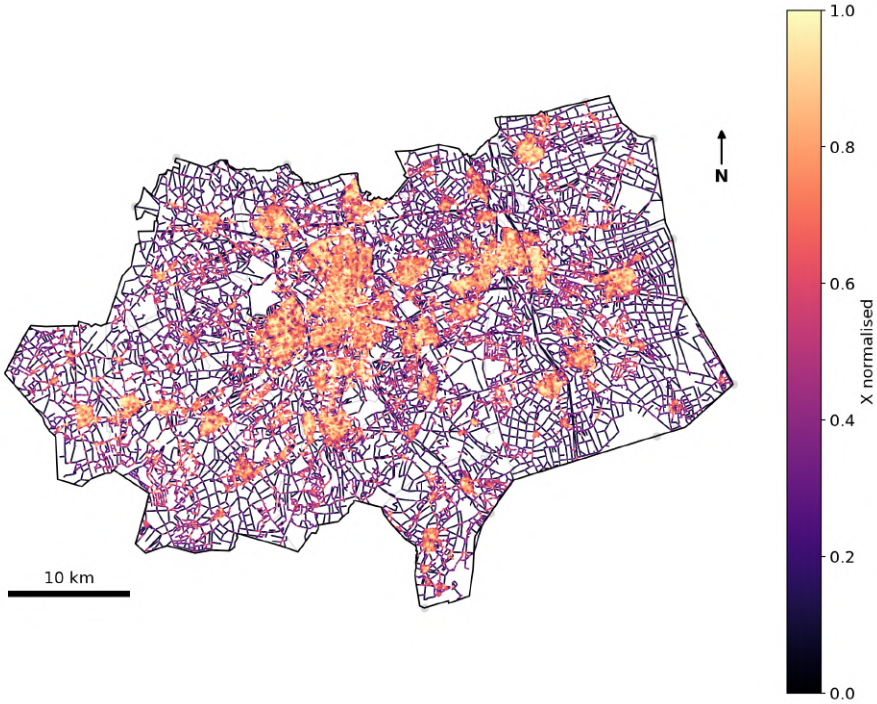

Figure A.1: Factor X: route-junction proximity.

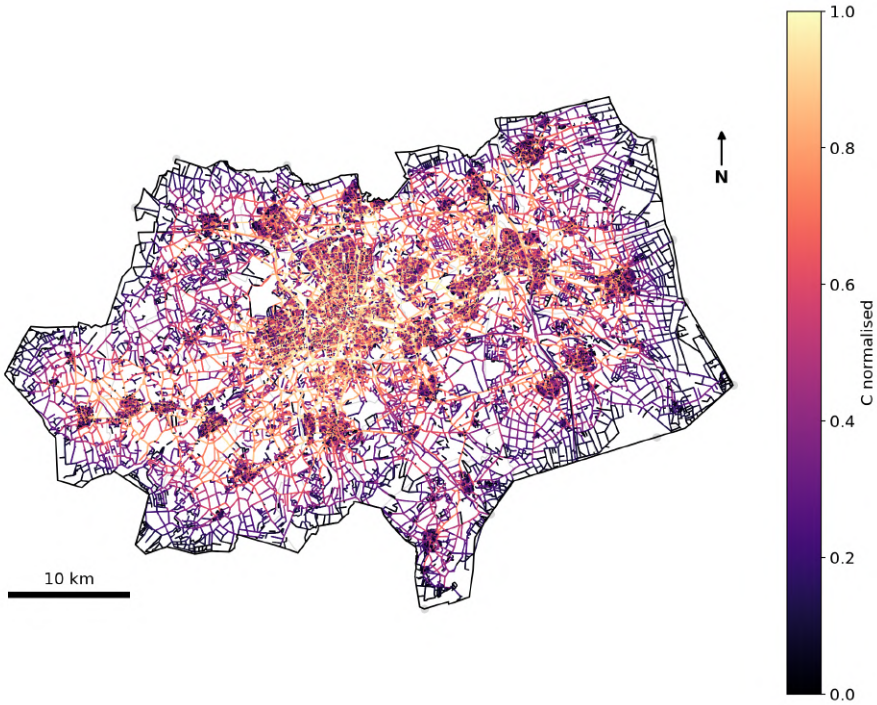

Figure A.2: Factor C: betweenness centrality.

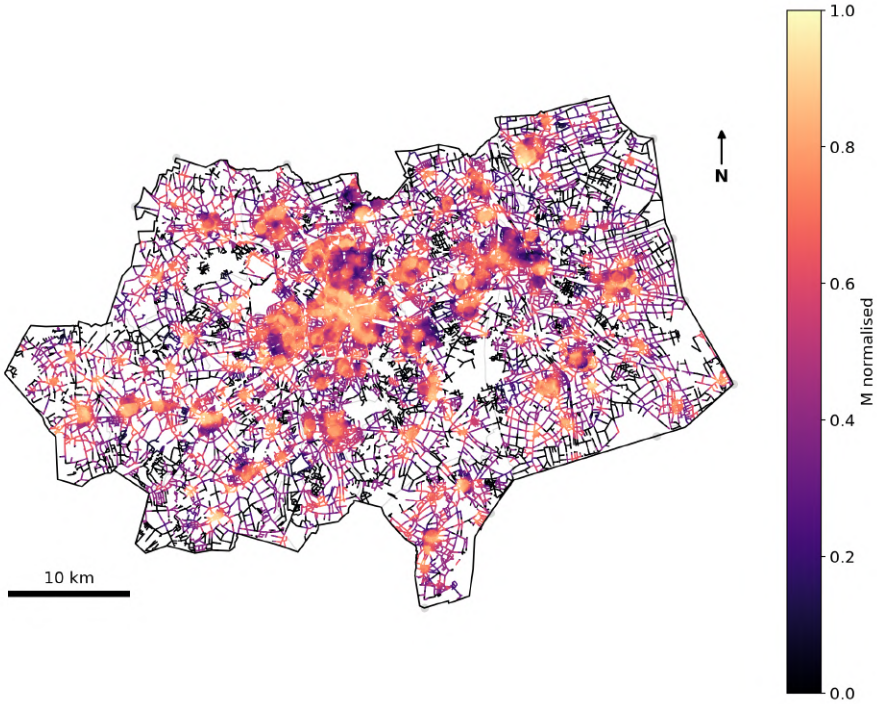

Figure A.3: Factor M: land-use mix.

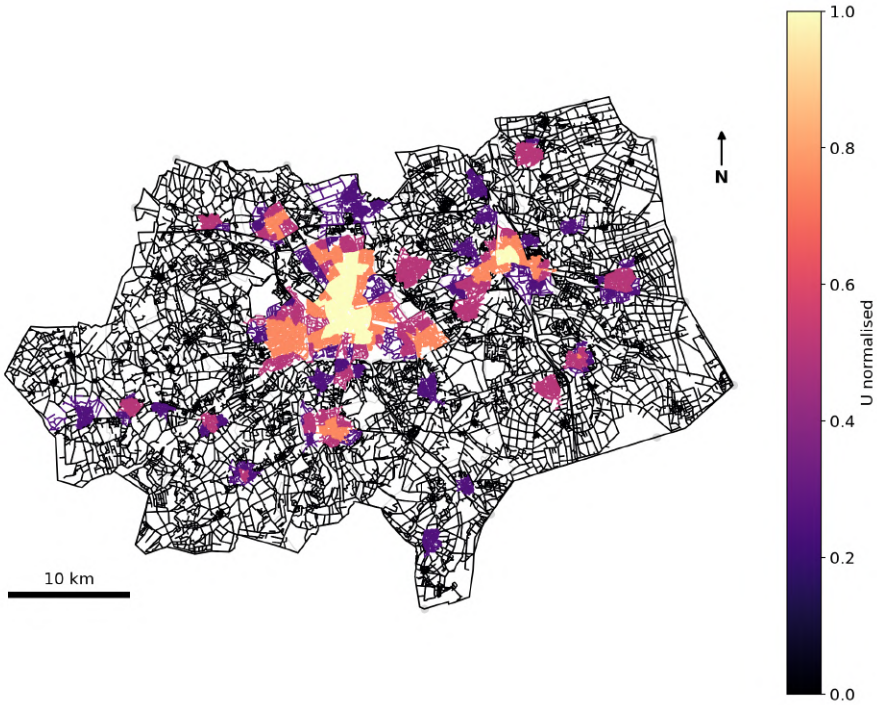

Figure A.4: Factor U: urban core (CBS urbanity).

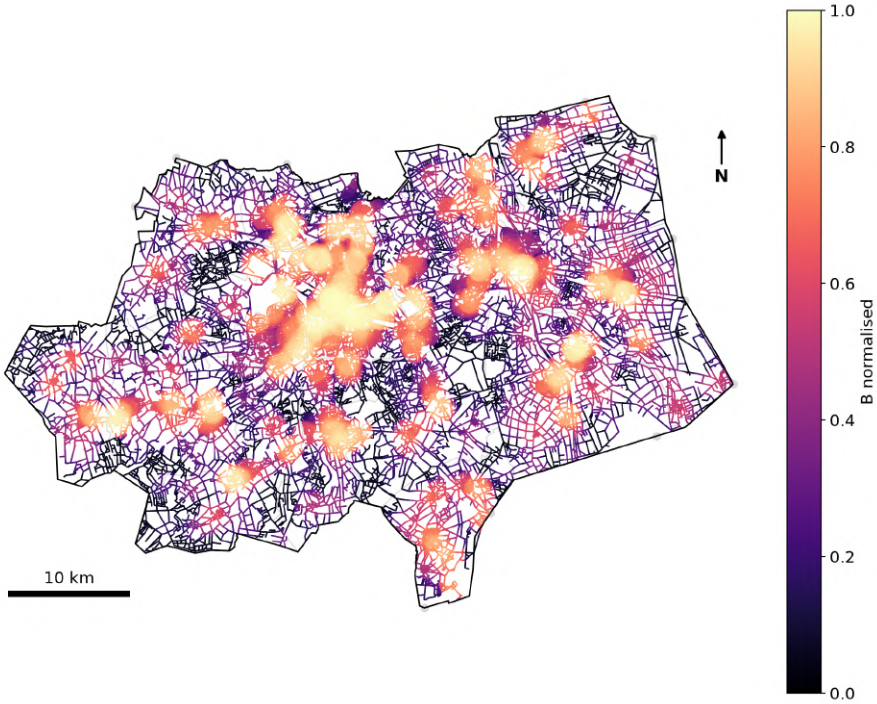

Figure A.5: Factor B: business density.

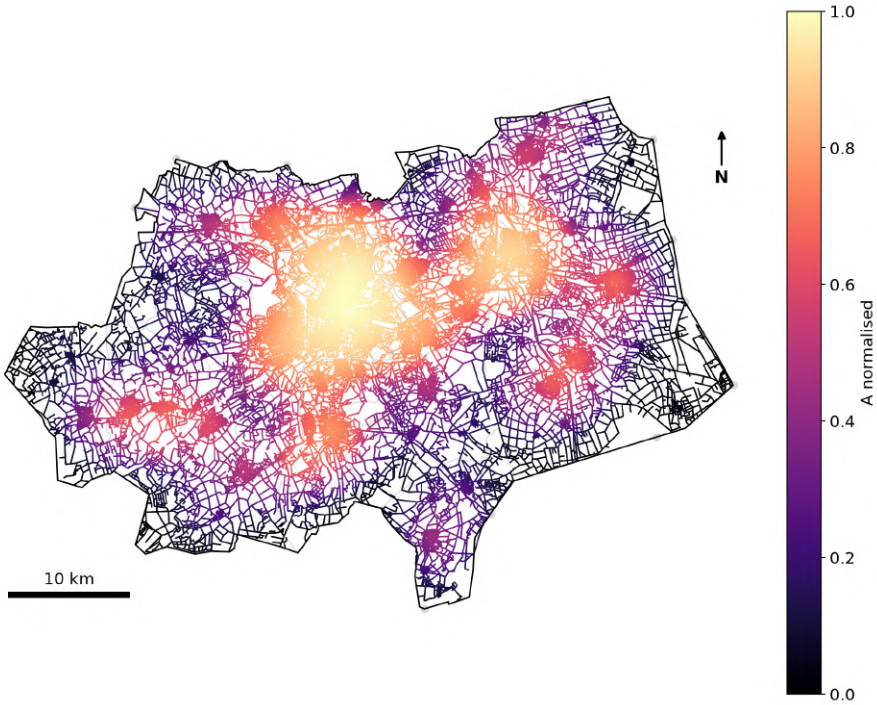

Figure A.6: Factor A: accessibility to opportunities.

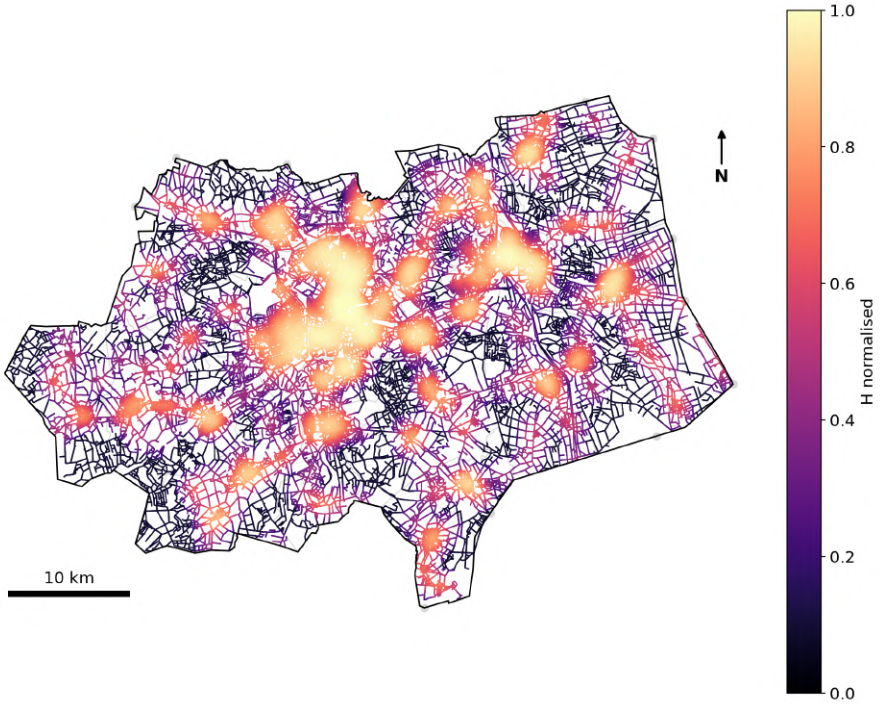

Figure A.7: Factor H: transit-hub proximity.

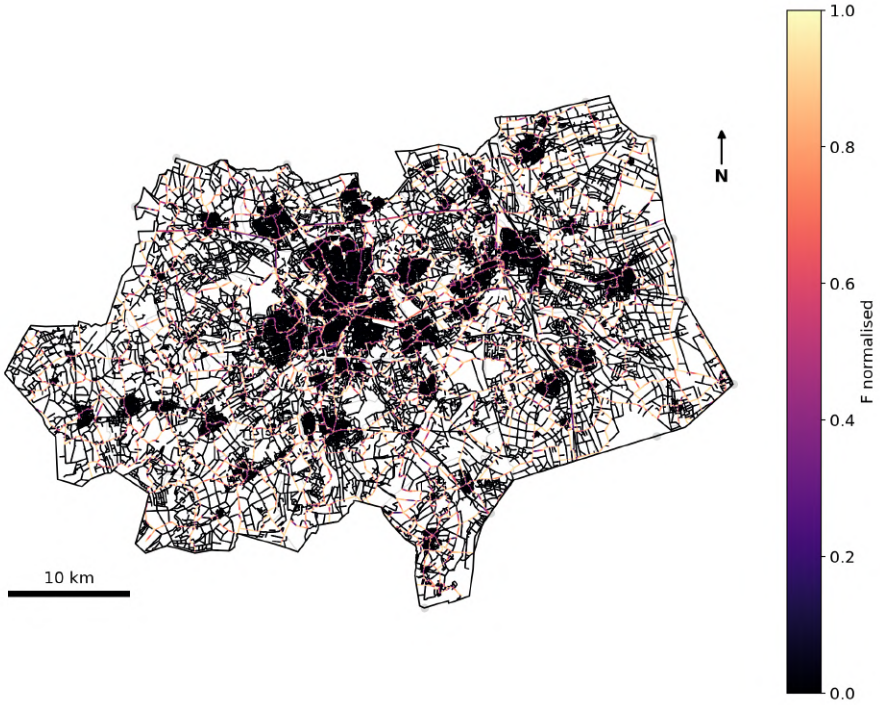

Figure A.8: Factor F: main cycle-network proximity.