CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

CCAT-NU-2025-5
December 2025

Making CAV Deployments Compatible with Complete Streets Objectives for Safe and Efficient Operation Phase II

Hani S. Mahmassani
Sharika J. Hegde
Meredith Raymer
Amirmohammad Khakpour

Northwestern

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Report No. CCAT-NU-2025-5

December 2025

Project Start Date: 6/1/2024

Project End Date: 12/22/2025

Making CAV Deployments Compatible with Complete Streets Objectives for Safe and Efficient Operation

Phase II

Hani S. Mahmassani

Professor

Sharika J. Hegde

Graduate Researcher

Meredith Raymer

Graduate Researcher

Amirmohammad Khakpour

Graduate Researcher

Northwestern University

DISCLAIMER

Funding for this research was provided by the Center for Connected and Automated Transportation under Grant No. 69A3551747105 of the U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology (OST-R), University Transportation Centers Program. The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation, University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Suggested APA Format Citation:
Mahmassani, H., Hegde, S., Raymer, M., & Khakpour, A. (2025). Making CAV deployments compatible with Complete Streets objectives for safe and efficient operation, Phase II (Final Project Report). Northwestern University, Center for Connected and Automated Transportation (CCAT).

Contacts
For more information:

Brett Johnson

Northwestern University
600 Foster St
Evanston, IL 60208

bretj@northwestern.edu
(847) 491-7287

CCAT

University of Michigan Transportation Research Institute
2901 Baxter Road
Ann Arbor, MI 48152

uumtri-ccat@umich.edu
(734) 763-2498

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Technical Report Documentation Page

1. Report No. CCAT-NU-2025-5	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Making CAV Deployments Compatible with Complete Streets Objectives for Safe and Efficient Operation – Phase II		5. Report Date December 2025 6. Performing Organization Code N/A
7. Author(s) Hani S. Mahmassani (orcid.org/0000-0002-8443-8928), Sharika J. Hegde (orcid.org/0000-0002-8822-3844), Meredith Raymer (orcid.org/0000-0003-2138-0842), Amirmohammad Khakpour (orcid.org/0009-0002-2441-5481)		8. Performing Organization Report No. N/A
9. Performing Organization Name and Address Northwestern University Transportation Center 600 Foster St. Evanston, IL, 60201		10. Work Unit No. 11. Contract or Grant No. Contract No. 69A3551747105
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE Washington, DC 20590		13. Type of Report and Period Covered Final Report (July 2024 – December 2025) 14. Sponsoring Agency Code OST-R
15. Supplementary Notes Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.		
16. Abstract		
17. Key Words Connected and Autonomous Vehicles, Complete Streets, Network Fundamental Diagram, Conflict management, Multimodal transportation	18. Distribution Statement No restrictions.	

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 70	22. Price

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Making CAV Deployments Compatible with Complete Streets Objectives for Safe and Efficient Operation

Phase II

Hani S. Mahmassani
https://orcid.org/0000-0002-8443-8928

Sharika J. Hegde
https://orcid.org/0000-0002-8822-3844

Meredith Raymer
https://orcid.org/0000-0003-2138-0842

Amirmohammad Khakpour
https://orcid.org/0009-0002-2441-5481

Northwestern University

Table of Contents

List of Figures

List of Tables

Abstract

Urban streets are increasingly complex environments where diverse transportation modes like cars, bicycles, pedestrians, micromobility, and emerging automated vehicles, interact within constrained spaces. Traditional design practices, guided by qualitative policies such as Complete Streets, often lack the quantitative foundation needed to evaluate trade-offs between efficiency, equity, and safety. This research introduces a data-driven framework to address these gaps by integrating macroscopic and microscopic modeling approaches for multimodal street design.

The study develops two major tools: a Network Fundamental Diagram (NFD)-based optimization framework and a multi-agent simulation platform. The optimization framework applies a two-step process: first, identifying flow-maximizing network attributes at the aggregate level, and second, allocating these attributes to specific links using a Genetic Algorithm (GA) under multiple ethical objectives, including utilitarian, sufficiency, accessibility gap, and maximin principles. The simulation platform models realistic interactions among vehicles, cyclists, pedestrians, and emerging modes, incorporating stochastic demand and behavioral rules to evaluate safety and efficiency. Conflict characterization combines Time-to-Collision (TTC) and Impact Severity metrics to assess both the likelihood and consequence of interactions, enabling nuanced safety evaluations.

Applications to real-world networks, including a case study of Evanston, IL and a corridor in Chicago, demonstrate the feasibility and effectiveness of these tools. Results show that optimized designs improve multimodal accessibility and reduce severe conflicts compared to existing layouts, while highlighting trade-offs between throughput, equity, and safety. Although no single design achieves all objectives simultaneously, the framework provides transparent evaluation and supports evidence-based decision-making. Policy priorities and community values remain essential in guiding the selection of final designs, ensuring that technological innovation aligns with human-centered goals.

1 INTRODUCTION

Urban streets are increasingly complex environments where diverse transportation modes like cars, bicycles, pedestrians, micromobility, and emerging automated vehicles, interact within constrained spaces. Traditional design practices, guided by qualitative policies such as Complete Streets, often lack the quantitative foundation needed to evaluate trade-offs between efficiency, equity, and safety. This research introduces a data-driven framework to address these gaps by integrating macroscopic and microscopic modeling approaches for multimodal street design.

This report addresses that need by developing and applying integrated modeling and optimization frameworks to make CAV deployments compatible with Complete Streets objectives. The research spans multiple scales of analysis: from macroscopic models that capture network-level flow dynamics to microscopic simulation platforms that represent agent-level behaviors and interactions. A unified approach combines Network Fundamental Diagram (NFD)-based optimization, multi-agent simulation, and conflict severity analysis to evaluate how design choices influence throughput, accessibility, and safety under realistic operating conditions. Ethical considerations such as sufficiency, equity, and maximin principles are incorporated into the optimization process to ensure that infrastructure decisions reflect not only technical efficiency but also social fairness.

By bridging the gap between policy-driven guidelines and quantitative modeling, this study provides urban planners and policymakers with actionable tools to design streets that balance technological innovation with human-centered values. The resulting frameworks enable transparent evaluation of trade-offs, support informed decision-making, and advance the vision of safe, efficient, and equitable urban mobility.

2 MICROSCOPIC BIKE NETWORK DESIGN PROBLEM

2.1 Introduction

This chapter outlines a methodology and results of a bike lane network design problem application to the real-world example of Evanston, IL. Building on the functional form of the NFD developed in Chapter 6, in this chapter a two-step optimization framework that integrates flow and accessibility maximization is developed In the first optimization, the functional form, which models flow as a function of network attributes (mode-exclusive area, shared area, and interaction plane) is used as the objective function, where the goal is to maximize the total flow of the network across all modes with the network attributes as decision variables. With these optimal network characteristics identified, the next step of the optimization aims to find the allocation of these characteristics to specific links. To do this, a genetic algorithm was implemented with the objective of maximizing accessibility and minimizing the difference between the predicted NFD curve for the network and the optimized NFD curve from the first step. The decision variables for this formulation are the set of links and their respective configurations. Alternative formulations that focus on meeting additional ethical considerations with the design are also included. These ethical frameworks consider the impacts on less advantaged groups, whether that be by mode used, income, or accessibility. A flowchart of this methodology, beginning with the functional form developed in Chapter 6, to the resulting network design is shown in Figure 2.1 .

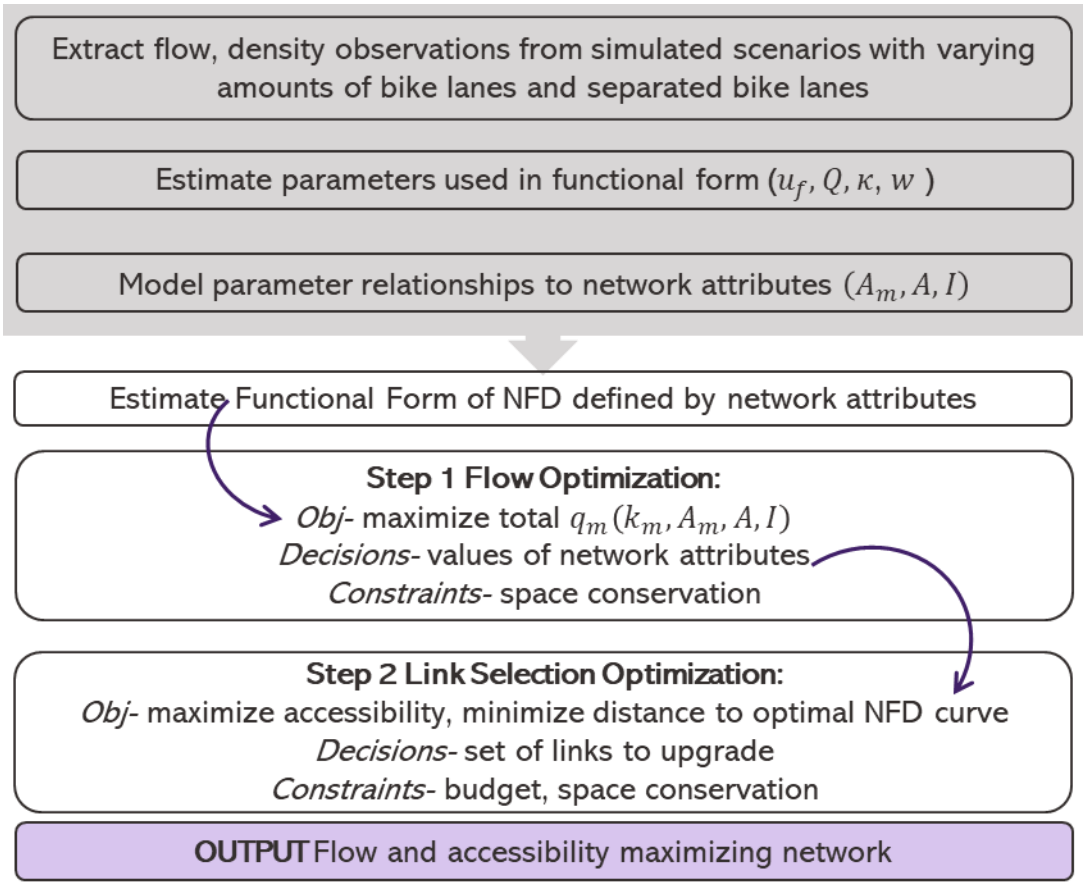

Figure 2.1 Process Flow of NFD led Bike Network Design Problem

2.2 Methodology

2.2.1 Step 1- Flow Optimization

The two step optimization process begins with the flow optimization formulation. Here, the functional form developed in Chapter 6 for $q_m(k_m,\ I,\ A_m,\ A\ )$ is maximized with decision variables $k_m,\ I,\ A_m,$ and A. The formulation of the optimization problem is outlined.

$$\max_{A_m, A,\ I | k_m} \sum_m q_m(k_m,\ I,\ A_m,\ A\ ) \quad \text{(Eq 2.1)}$$

s.t.

$$A + \sum\nolimits_m A_m < A_{max} \quad \text{(Eq 2.2)}$$

$$0 < k_{crit,m} < \beta_{7m} + \beta_{8m} A_m + \beta_{9m} I \quad \text{(Eq 2.3)}$$

$$I,\ A_m,\ A\ \geq 0 \quad \text{(Eq 2.4)}$$

where k_m is the density of mode m, I is the interaction plane, A_m is the modal area of mode m, A is the shared area, and A_{max} is the total existing area of the network.

The inner objective finds the critical density k_m where the flow of the given function is at a maximum given the values of the other three variables. $I,\ A_m,\ A,$ are the outer optimization's decision variables. Across the inner and outer optimization outlined in Equation 2.1, flow is maximized across all modes; in this example, bikes and cars are the modes considered. The first constraint, Equation 2.2, conserves the existing area of the network- no additional area can be added to the network to increase flow. The second constraint, Equation 2.3, keeps the estimation of critical density feasible, less than the value of jam density κ developed in the estimation of the NFD form. The last constraint, Equation 2.4, ensures that the values of the network attribute decision variables $I,\ a_m,$ and A are non-negative.

2.2.2 Step 1- Solution Method

The Step 1 formulation was solved using the Python package *Pyomo* and the *IPOPT* solver to address a nonlinear, constrained optimization problem. *IPOPT* is a large-scale, interior-point optimization solver designed to handle smooth, nonlinear objectives with both equality and inequality constraints. *IPOPT* guarantees convergence to a locally optimal solution under smoothness assumptions, but global optimality cannot be assured due to the non-convex nature of the problem.

2.2.3 Step 2- Space Allocation Optimization

Step 1 results in flow maximizing values of I, a_m, and A for the whole network. The second step of the framework aims to allocate these values to specific links in the network in order to meet these aggregate optimal values, resulting in an optimal NFD. Given a set of links and outcomes, $x_{l,o}$ is the decision variable to make a link l have an outcome type o where o represents a possible treatment and a corresponding width. The possible treatments are shown in Figure 2.2 .

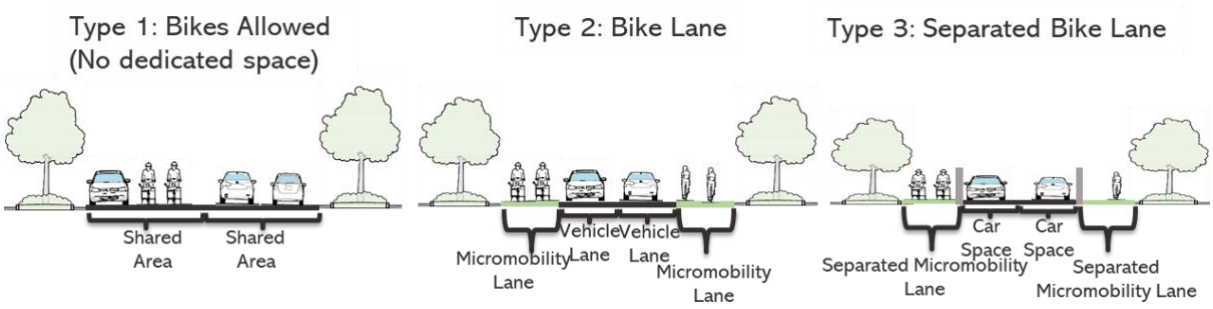

Figure 2.2 Link Treatment Types

Accessibility is defined by $z_{i,\,m|20}(x_{l,o})$ which is the number of zonal centroids accessible from the centroid of zone i by mode m within a 20-minute threshold for the network configuration $x_{l,o}$. The minimization portion of the objective aims to reduce the area between the optimal NFD curve found in Step 1 and the resulting NFD curve for a configuration $x_{l,o}$. The formulation for Step 2 is shown.

$$\max \sum_{M} \sum_{I} z_{i,\,m|20}(x_{l,o}) - \sum_{m} \left(\int_{0}^{\kappa} \left[q_m^*(k_m^*, I^*, A_m^*, A^*) - q_m(k_m, I, A_m, A) \right] dk \right) \quad \text{(Eq 2.5)}$$

s.t.

$$\sum_{o \in O} x_{l,o} = 1 \quad \forall l \in L \quad \text{(Eq 2.6)}$$

$$\sum_{l \in L} \sum_{o \in O} x_{l,o} C_{l,o} \leq B \quad \text{(Eq 2.7)}$$

$$\sum_{l \in L} \sum_{o \in O} \left(A(l,o) + \sum_{m} A_m(l,o) \right) < A_{max} \quad \text{(Eq 2.8)}$$

$$\sum_{o} w_{l,o}\, x_{l,o} \leq W_l \quad \text{(Eq 2.9)}$$

$$\sum_{o} w_{l,o}\, x_{l,o} \leq W_l \qquad \text{(Eq 2.10)}$$

where $C_{l,o}$ is the construction cost of a configuration $x_{l,o}$, W_l is the total width of the individual link l.

The objective function, Equation 2.5, maximizes the total accessibility within a 20-minute threshold across modes m and zones i while minimizing the area between NFD curves through the integral difference. The first constraint, Equation 2.6, ensures that each link only has one outcome o. Equation 2.7 is a budget constraint where each network configuration has a cost of implementation $C_{l,o}$. Equation 2.8 imposes a total area constraint which conserves the total network area. Equation 2.9 further specifies this to each link, which cannot gain additional width. Equation 2.9 forces the total values of $I,\ a_m,\ A,$ and $Z_{i,\,m|20}$ to be non-negative.

2.2.4 Step 2- Modified Space Allocation Optimization using Area Constraints

An alternative formulation of the second-step optimization is presented for cases where the actual values obtained in Step 1 are unavailable or infeasible to implement. In such scenarios, the allocation of interaction plane length *(I),* mode-specific areas *(a_m),* and total area *(A)* may instead be determined based on local priorities and needs. In this version of the formulation, values of *I, a_m,* and *A* are treated as fixed constraints, and accessibility is maximized directly as opposed to minimizing the area between the predicted and optimal NFDs.

$$max \sum_{M} \sum_{I} z_{i,\,m|20}(x_{l,o}) \qquad \text{(Eq 2.11)}$$

s.t.

$$\sum_{o \in O} x_{l,o} = 1 \ \ \forall l \in L \qquad \text{(Eq 2.12)}$$

$$\sum_{l \in L} \sum_{o \in O} x_{l,\,o} C_{l,o} \ \leq B \qquad \text{(Eq 2.13)}$$

$$\sum_{l \in L} \sum_{o \in O} (A(l,o) + \sum_{m} A_m(l,o)) < A_{max} \qquad \text{(Eq 2.14)}$$

$$\sum_{o} w_{l,o}\, x_{l,o} \leq W_l \qquad \text{(Eq 2.15)}$$

$$I,\ A_m,\ A, Z_{i,\,m|20} \geq 0 \qquad \text{(Eq 2.16)}$$

$$\sum_{l \in L} \sum_{o \in O} A(l,o) \geq A^* - A^* 5\% \quad \text{(Eq 2.17)}$$

$$\sum_{l \in L} \sum_{o \in O} A(l,o) \leq A^* + A^* 5\% \quad \text{(Eq 2.18)}$$

$$\sum_{l \in L} \sum_{o \in O} A_m(l,o) \geq A_m^* - A_m^* 5\% \quad \text{(Eq 2.19)}$$

$$\sum_{l \in L} \sum_{o \in O} A_m(l,o) \leq A_m^* + A_m^* 5\% \quad \text{(Eq 2.20)}$$

$$\sum_{l \in L} \sum_{o \in O} I(l,o) \geq I^* - I^* 5\% \quad \text{(Eq 2.21)}$$

$$\sum_{l \in L} \sum_{o \in O} I(l,o) \leq I^* + I^* 5\% \quad \text{(Eq 2.22)}$$

Equations 2.12 - 2.16 reflect constraints on total area, budget, link width, and non-negativity as outlined in the initial formulation. Equations 2.17 and 2.18 2.182.18restrict the shared area A to be within a ±5% buffer of the optimal value A^* identified in Step 1. Equations 2.19 and 2.20 impose the same restriction for modal area, A_m^*. Equations 2.21 and 2.22 aim to meet the optimal interaction plane value I^* of Step 1 within a ±5% buffer.

2.2.5 Step 2- Space Allocation Optimization with Additional Ethical Formulations

Both the NFD-based and optimal value-based formulations fall under a utilitarian framework of ethics where the objective is maximizing total accessibility. By modifying the objective function and introducing additional constraints aligned with ethical theories, this study adapts applications of ethical theory, originally developed for transit network design, to the context of bike network design [1]. The initial formulation covers a utilitarian approach. Other principles considered are the sufficiency principle sets a minimum threshold of accessibility to ensure no resident is excluded from basic needs, while the accessibility gap limits inequality by regulating the maximum allowable difference in accessibility across residents. The maximin formulation prioritizes maximizing accessibility for the least advantaged.

2.2.5.1 Sufficiency Criteria

The first additional constraint is based on the ethics framework of sufficiency, which ensures that all areas meet some minimum threshold of accessibility. Sufficiency is implemented by the addition of a constraint to either formulation.

$$\sum_{M} \sum_{I} z_{i,\, m|45}(x_{l,o}) \geq S \quad \text{(Eq 2.23)}$$

Equation 2.23 ensures zonal accessibility within a 45-minute threshold should be greater than S which is the threshold number of zonal accessibility to be reached.

2.2.5.2 Accessibility Gap Criteria

The accessibility gap criterion aims to limit the difference in accessibility between the worst-off and best-off areas. Worst-off and best-off can be defined in many ways; the difference between zones of varying income, transportation modes, and accessibility are considered. The constraints included for each of these considerations are shown in Equations 2.24, 2.25, and 2.26, respectively.

$$\sum_{M} Z_{LI,\,m|20} - \sum_{M} Z_{HI,\,m|20} \geq 0 \qquad \text{(Eq 2.24)}$$

$$\sum_{I} Z_{i,\,B|20} - \sum_{I} Z_{i,\,C|20} \geq 0 \qquad \text{(Eq 2.25)}$$

$$\sum_{M} \min\left(Z_{i,\,m|20}\right) - \sum_{M} \max\left(Z_{i,\,m|20}\right) \geq 0 \qquad \text{(Eq 2.26)}$$

Equations 2.24 ensures that the lowest income zone, *LI,* receives equal or better total accessibility than the highest income zone, *HI*. Equation 2.25 ensures that total accessibility by bikes is at least as great as that by cars across all zones. Equation 2.26 ensures there is no difference between the highest accessibility zone and the lowest accessibility zone. Each constraint is added independently to the formulation giving three possible accessibility gap results based on income, mode, and accessibility.

2.2.5.3 Maximin Criteria

The final ethical criterion considered is the maximin framework, which seeks to maximize the outcomes for the least advantaged group. With the maximin formulation, the first term of Equation 2.5 is changed to reflect disadvantage on the basis of income, mode, and accessibility.

$$max \sum_{M} \sum_{I} Z_{LI,\,m|20}\left(x_{l,o}\right) \qquad \text{(Eq 2.27)}$$

$$max \sum_{I} Z_{i,\,B|20}\left(x_{l,o}\right) \qquad \text{(Eq 2.28)}$$

$$max \sum_{M} \min\left(Z_{i,\,m|20}\right) \qquad \text{(Eq 2.29)}$$

Equation 2.27 maximizes the access of the lowest income zone, *LI*. Equation 2.28 maximizes accessibility by bikes. Equation 2.29 maximizes the accessibility of the lowest accessibility zone. When using the Modified Space Allocation Optimization using Area Constraints, the maximin criteria objectives may be directly swapped with Equation 2.11. When using the Space Allocation Optimization with NFD Area Minimization, however, the objective function for each maximin criteria implementation would be as follows.

$$max \sum_{M} \sum_{I} Z_{LI,\, m|20}(x_{l,o}) - \sum_{m} (\int_{0}^{\kappa} q_m^*(k_m^*, I^*, A_m^*, A^*) - q_m(k_m, I, A_m, A\,)\, dk) \quad \text{(Eq 2.30)}$$

$$max \sum_{I} Z_{i,\, B|20}(x_{l,o}) - \sum_{m} (\int_{0}^{\kappa} q_m^*(k_m^*, I^*, A_m^*, A^*) - q_m(k_m, I, A_m, A\,)\, dk) \quad \text{(Eq 2.31)}$$

$$max \sum_{M} min\,(Z_{i,\, m|20}) - \sum_{m} (\int_{0}^{\kappa} q_m^*(k_m^*, I^*, A_m^*, A^*) - q_m(k_m, I, A_m, A\,)\, dk) \quad \text{(Eq 2.32)}$$

In Equations 2.30, 2.31, and 2.32 the minimization of the maximin between the optimal and predicted NFD curves is included. These objective functions are each swapped with the objective outlined in the formulation of Equations 2.5 - 2.10 where the constraints are shown by Equations 2.6 - 2.10 remain.

2.2.6 Step 2- Solution Algorithm

These formulations are implemented and solved using a Genetic Algorithm (GA) with Python's DEAP package. In this context, an individual represents a specific configuration of the network, where each edge is assigned a value for $x_{l,o}$ denoting the type and width of bike infrastructure on that link. Each individual corresponds to a complete SUMO network reflecting this configuration. To evaluate network performance, mode-specific trips are generated from the centroid of each zone and routed to all other zones using the SUMO tool, *duarouter*, resulting in trip lengths for each origin-destination pair and mode. For cyclists, this trip length is additionally updated with the marginal rate of substitution of the cyclist stress index developed in the literature [2]. Links with higher speed limits and less dedicated bike infrastructure increase the stress index of cyclists, resulting in effective trip lengths of up to 100% longer, as shown in Table 2.1.

Table 2.1 Marginal Rate of Substitution by Bike Lane Type

Speed Limit	Type 1	Type 2	Type 3
25 mph	20%	10%	5%
30 mph	40%	20%	10%
35 mph	100%	50%	25%

Using the defined free flow speed, the distances of these routes are converted to travel times between all zones by each mode, resulting in the count of OD pairs accessible within defined time thresholds.

Construction costs are included for both the addition of bike lanes and the removal of bike lanes that exist in the current network. The initial budget is set at $10 million, similar to other projects that have been completed in the city. The values considered are estimates from a recent construction proposal and are shown in Table 2.1 [3].

Table 2.2 Construction Costs by Infrastructure Element

Item	Units	Unit Cost
Painting Bike Lane	sq.ft.	$14
Installing Separation (curb)	ft.	$30

Paint removal	sq.ft.	$10
Separation (curb) removal	ft.	$9
Repaving	sq. yd.	$30

Within the Genetic Algorithm (GA), several functions guide the search for optimal solutions. Because random mutations often produce disjointed or unrealistic network configurations, a custom operator was used to develop more geographically contiguous solutions. The function biases mutation decisions based on the configuration of adjacent edges. During mutation, the function iterates through each edge in the individual solution. For a given edge, it retrieves the types of its neighbors and identifies the most frequently occurring type. If the current edge differs from this type and a mutation probability threshold is met, the edge is reassigned to the majority type among its neighbors. The crossover function combines parent solutions by swapping edge types and occasionally re-randomizing both offspring. Additionally, a custom repair function was developed to eliminate small, disconnected links of Type 2 or 3. The function ensures that any group of contiguous edges assigned to the same type meets a minimum cluster size. This repair mechanism is applied after mutation and crossover to enforce a minimum spatial structure in solutions and reduce fragmented configurations. Finally, a feasibility check ensures that individuals remain within the construction budget before being admitted into the next generation. The construction cost of each configuration is calculated based on the cumulative width and length of added or modified bike lanes. Individual solutions that exceed the budget are either excluded or penalized, depending on their performance trade-offs. This soft constraint allows the budget to be exceeded when substantial gains in accessibility or NFD alignment justify the additional cost. Tournament selection is used to identify the fittest individuals from each generation and guide the selection. In a *k*-way tournament, *k* candidates are randomly chosen, and the one with the highest fitness is selected to proceed.

This process is repeated to form the next generation until convergence of solutions is reached. To determine when to stop the evolutionary process, convergence is measured using a tolerance-based stopping criterion. Convergence is measured by tracking changes in the best fitness value across generations. If the improvement is smaller than a set tolerance and number of consecutive generations, the algorithm is considered to have converged and stops. This approach ensures that the algorithm stops when improvements become negligible.

2.3 Results and Discussion

2.3.1 Step 1 Optimization Results

The first optimization step finds the optimal aggregate values of network attributes, shared area, exclusive modal area, and interaction plane as described in the methodology. This nonlinear, constrained optimization problem was implemented using *Pyomo* with the *IPOPT* solver. The implementation of this formulation resulted in solutions that converged after 51 iterations, as shown in Figure 2.3 .

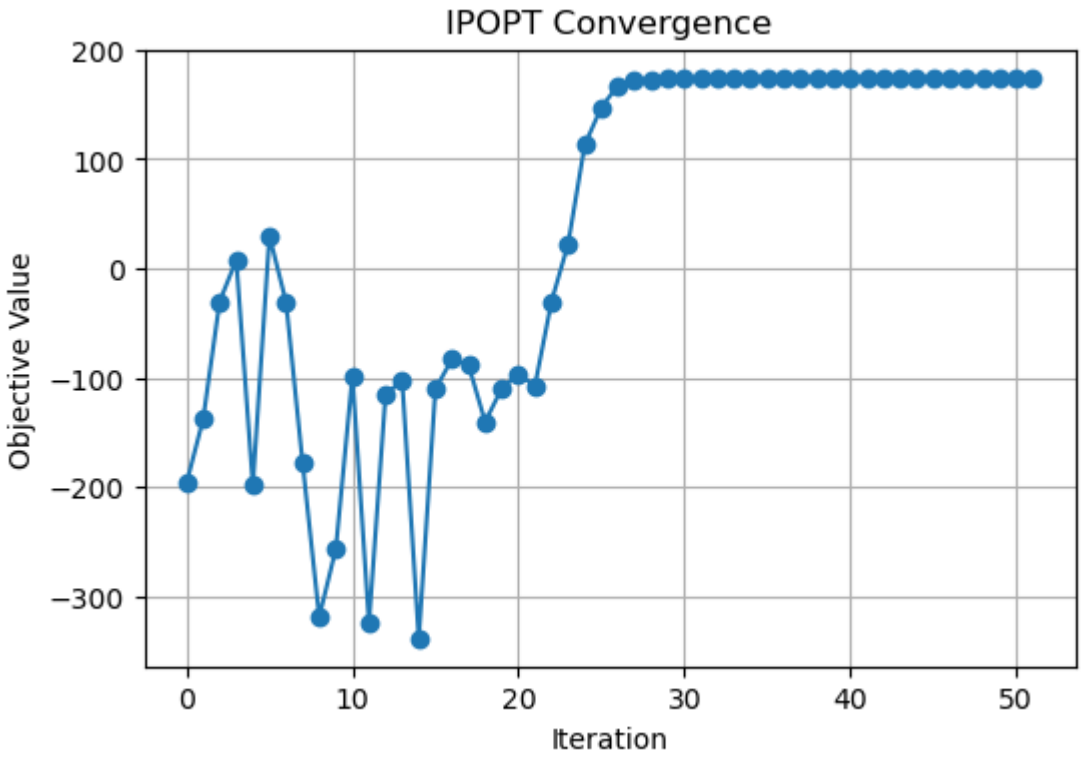

Figure 2.3 Convergence of Step 1 Optimization

In *IPOPT*, convergence indicates that the first-order optimality conditions have been met within specified tolerances, thus a locally optimal solution was obtained for the decision variables. The resulting values of these network attribute decision variables (shared area, car area, bike area, and interactions) that optimize the flow function q, as well as the locally optimal objective value, are shown in

Table 2.3.

Table 2.3 Optimal Values of Decision Variables and Objective from Step 1

Variable	Units	Optimal Value
Total Flow	veh/m/min	174.53
Car Area	km^2	0.91
Bike Area	km^2	0.88
Shared Area	km^2	0.16
Interaction Plane	km	53.00
Critical Density Car	veh/m^2	0.47
Critical Density Bike	veh/m^2	0.98

The optimal values include a car area and a bike area of similar values, 0.91 km^2 and 0.88 km^2, respectively. The shared area of 0.16 km^2 reflects a limited integration of shared space, which has a negative relationship with capacity. This shared area is also associated with a smaller interaction plane value of 53.00 km, which was previously identified as having a negative relationship with capacity in the functional form for bikes. The resulting optimal curves are shown in Figure 2.4 for both cars and bikes.

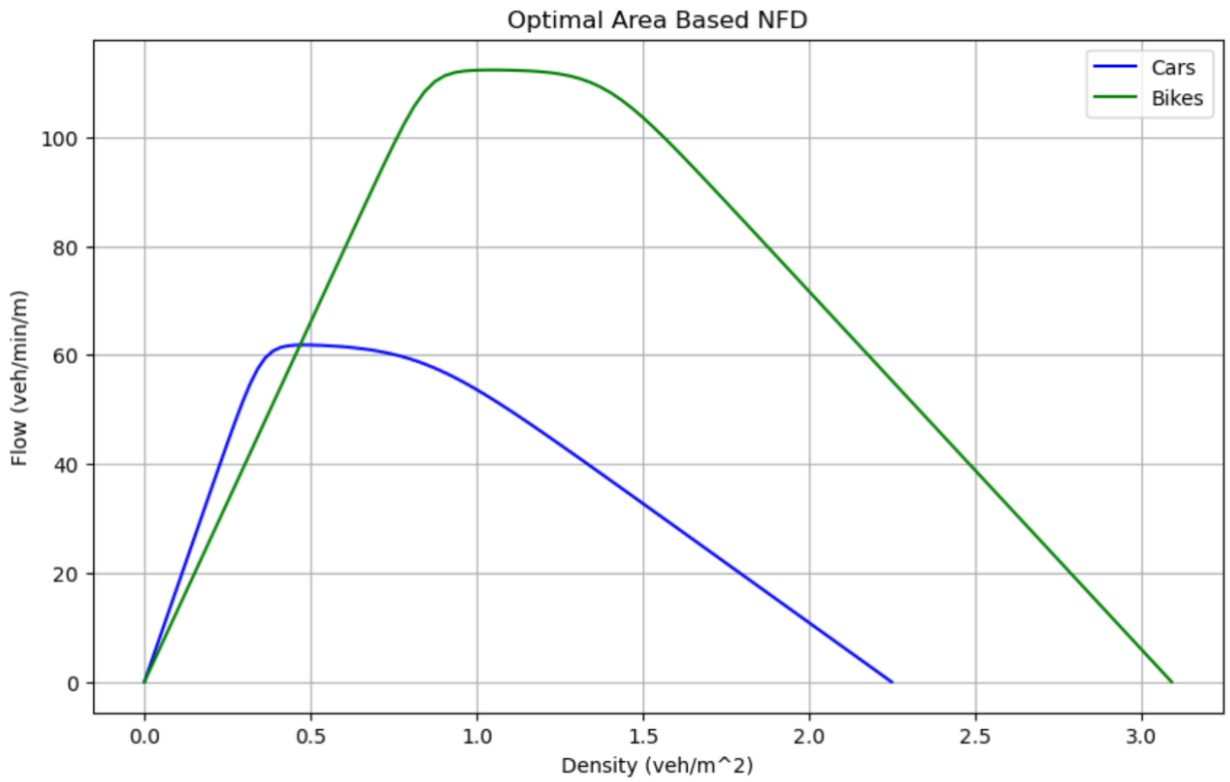

Figure 2.4 Optimal Form of NFD for Evanston Network

The micromobility curve (green) demonstrates a higher overall capacity, peaking at approximately 110 veh/min/m near a density of 1 veh/m². The car mode curve (blue) reaches a lower peak flow of about 60 veh/min/m at a density just below 0.5 veh/m². This also reflects the efficiency of micromobility vehicles, which contribute more to the total flow being maximized. Though cars and bikes are allocated almost equal mode-exclusive areas in the decision variables, bikes achieve almost double the flow with higher densities.

2.3.2 Step 2 Optimization Results- Utilitarian Formulation

This optimal curve is then implemented in the GA framework with the formulation described in Equations 2.5 - 2.10, which aims to generate a network configuration that both minimizes the area difference between its estimated NFD and the optimal curve and maximizes accessibility. This section presents the results and analysis of the network solution, which is based on utilitarian allocation. To evaluate the performance of the Genetic Algorithm across different formulations, convergence behavior and population diversity are tracked over

generations. Figure 2.5 illustrates the progression of best and average fitness values as well as the diversity of solutions to validate the convergence towards near-optimal solutions.

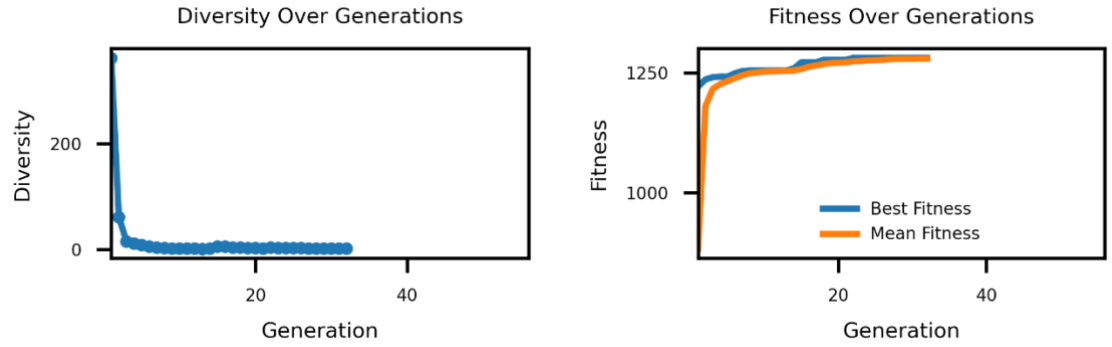

Figure 2.5 Convergence of Utilitarian Solution

Both best and mean fitness demonstrate improvement within the first 10 generations, with convergence observed shortly thereafter. The diversity of the population over generations measures the spread of solutions. The decline in diversity in early generations reflects strong selection pressure, followed by stabilization near zero, suggesting convergence to a narrow region of the solution space.

The GA seeks to minimize the area between a solution network's NFD and the optimal NFD found in Step 1; each curve and the corresponding difference in area are shown in Figure 2.6.

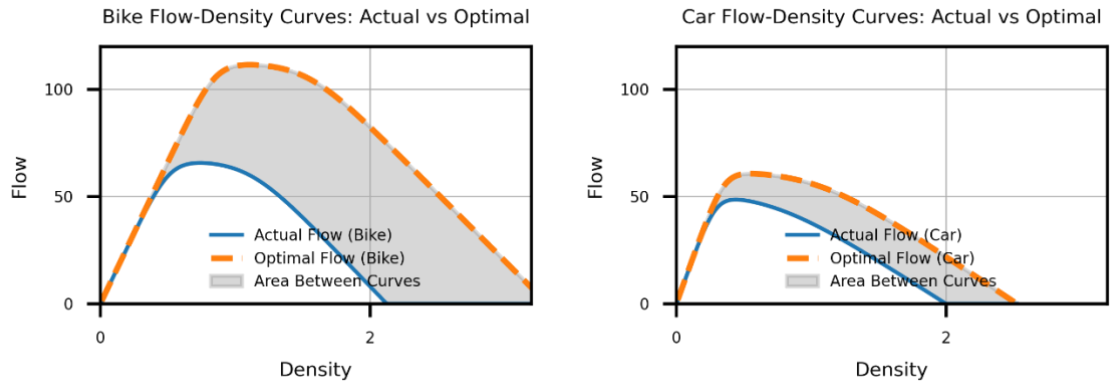

Figure 2.6 NFD and Optimal NFD of Utilitarian Network Design

Across the simulated scenarios and the Step 1 optimization, the efficiency of micromobility vehicles has been highlighted. However, when considering the area between the curves, the car NFD is able to achieve closer-to-optimal performance. This discrepancy arises from the solution not reaching the specific micromobility-exclusive area values that yield optimal flow, likely due to budget constraints. The actual network with bike lane link type selections can be seen in Figure 2.7.

Figure 2.7 Utilitarian Design of Evanston Bike Network with Type 2 and Type 3 lanes

To best visualize the locations of bike lanes, the network is shown twice, first with Type 2 lanes highlighted and then Type 3 lanes highlighted. The optimized network demonstrates a widespread and fairly uniform selection of Type 3 links (teal), particularly concentrated along major corridors and key north-south and east-west connectors. Type 2 (orange) lanes are added to local streets and peripheral areas. Overall, the spatial pattern indicates a strategy of concentrating high-quality infrastructure along primary routes aligning with a utilitarian goal of maximizing system-wide accessibility and flow. While the distribution of the network links is valuable, understanding the effectiveness of the design requires examining the underlying performance metrics, which are quantified in Table 2.4.

Table 2.4 Metrics of Utilitarian GA Solution

Metric	Units	Value	Target/Upper Bound
Car 20-min Accessibility	# trips	2109	3306
Bike 20-min Accessibility	# trips	840	3306
Total 20-min Accessibility	# trips	2949	6612
Construction Cost	$	9,974,335	10,000,000
Car Area	km^2	0.43	0.91
Shared Area	km^2	1.13	0.16
Bike Area	km^2	0.40	0.88
Interaction Plane	km	228.02	53.00
Area between Car NFD (optimal and actual)	veh^2/m^3*min	45.61	0.00
Area between Bike NFD (optimal and actual)	veh^2/m^3*min	133.18	0.00

In terms of 20-minute accessibility, the algorithm achieves 2,949 connected zones for both car and bike modes, aligning with the objective of maximizing overall accessibility. However, this underperforms relative to the target, reaching only 2,109 car trips and 840 bike trips. Car and bike areas fall below the desired 0.91 km² and 0.88 km², respectively, while shared area remains over-allocated at 1.13 km² compared to the optimal 0.16 km². Both car and bike networks exhibit divergence from the ideal curves (45.61 and 133.18 veh²/m³/min, respectively).

2.3.3 Step 2 Optimization Results- Additional Ethical Formulations

In this section, comparisons of outcomes under alternative ethical frameworks are shown. These include sufficiency- and priority-based approaches, which emphasize equitable access rather than total network efficiency. Figure 2.8 presents the convergence behavior for each fairness-oriented objective tested in the Genetic Algorithm framework.

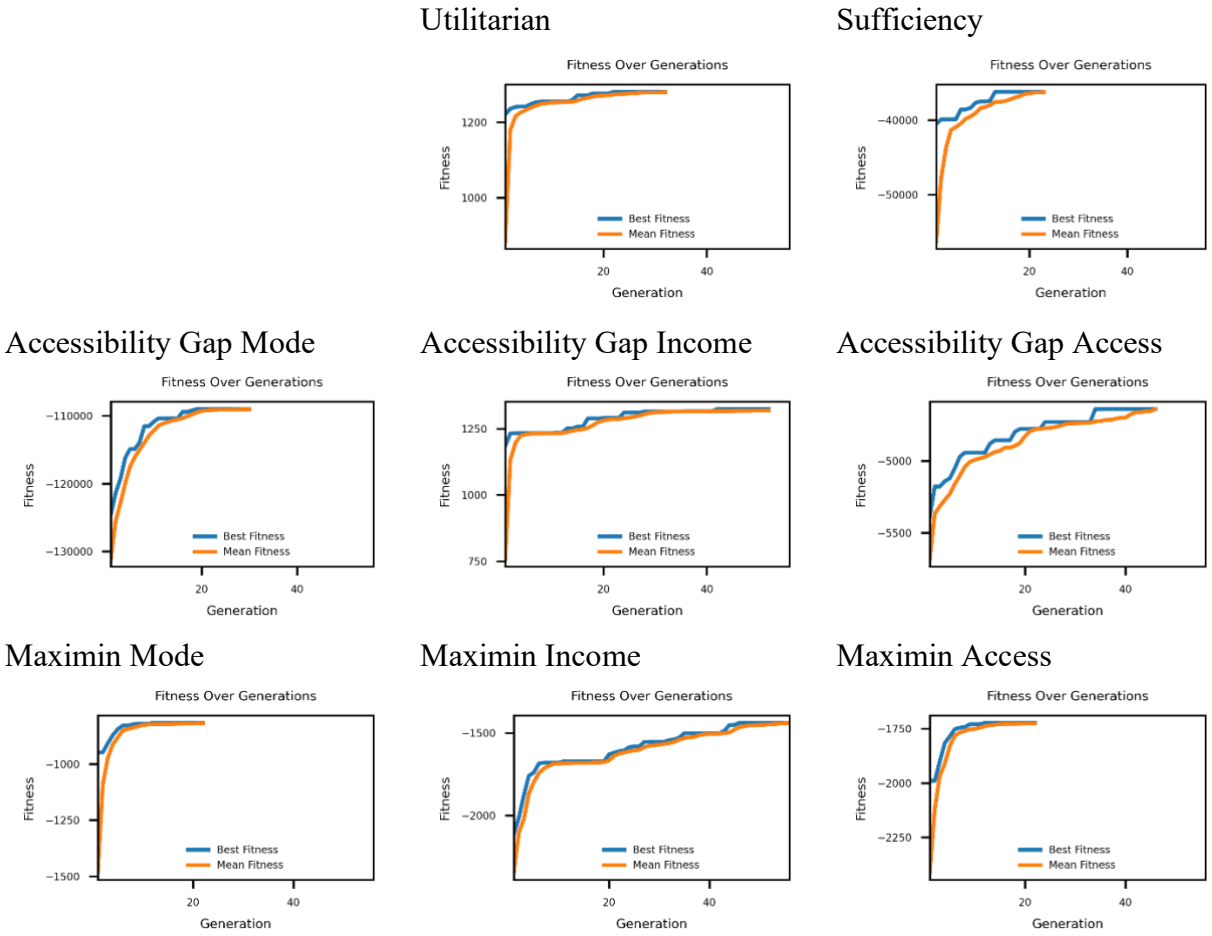

Figure 2.8 Convergence of GA Over Fairness Objectives

Across all objectives, the algorithm exhibits convergence within the first few generations. For most cases, both best and mean fitness improve steadily before plateauing, indicating that the population quickly evolves toward high-quality solutions. Notably, objectives like *Accessibility Gap Income* and *Accessibility Gap Access* show continued incremental gains across nearly 50

generations, suggesting a slower and more prolonged search process. In contrast, objectives such as *Utilitarian*, *Maximin Mode*, and *Maximin Access* converge rapidly, often stabilizing within the first 10–15 generations. These differences indicate varying levels of search complexity: some objectives reach high-performing solutions quickly, while others require sustained optimization over a longer period to reach convergence. Figure 2.9 shows diversity, measured as the standard deviation across individuals in each generation.

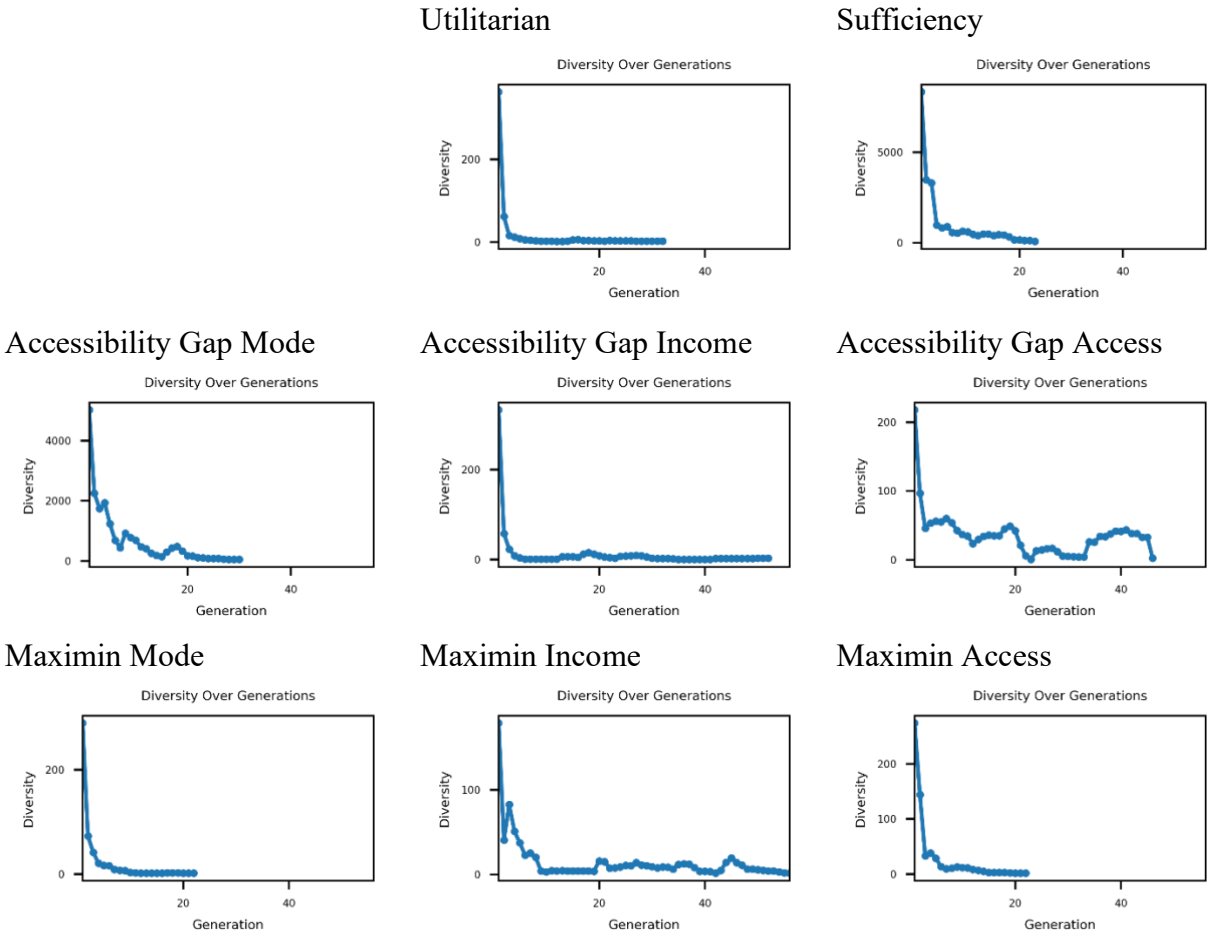

Figure 2.9 Diversity over Generations of GA

The diversity plots generally show a sharp drop in the early generations, which is expected as the algorithm quickly exploits promising regions of the search space. An exception to this trend is observed in *Accessibility Gap Access,* where diversity remains volatile. This reflects trade-offs within the solutions where improving accessibility in one low-access area may inadvertently reduce accessibility in another.

To evaluate how well each network aligns with the optimal NFD found in Step 1, the estimated NFD curves of networks generated by the GA are compared to the optimal curve. This optimal curve represents the theoretical upper bound of flow performance achievable. The plots in

Existing Utilitarian Sufficiency

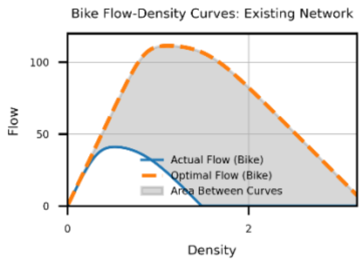

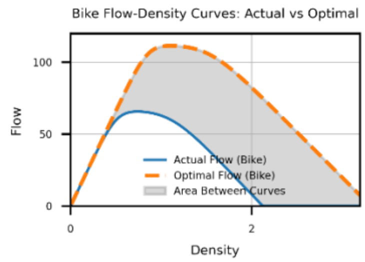

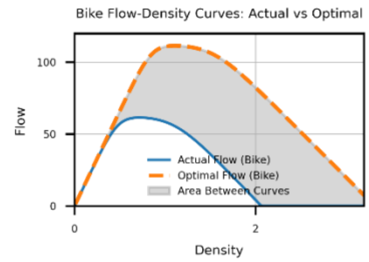

Accessibility Gap Mode

Accessibility Gap Income

Accessibility Gap Access

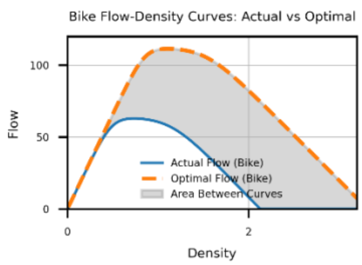

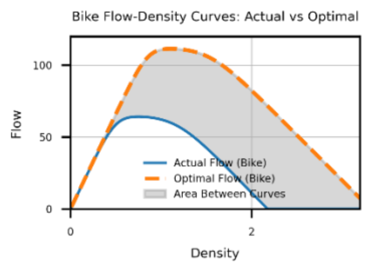

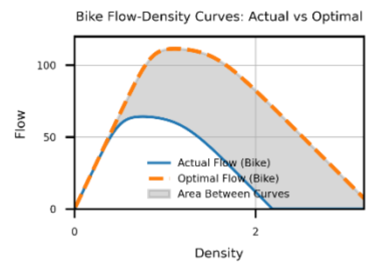

Maximin Mode

Maximin Income

Maximin Access

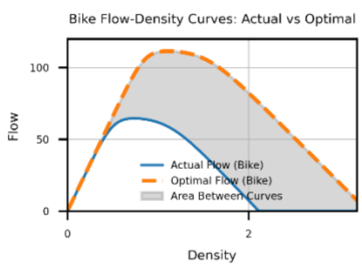

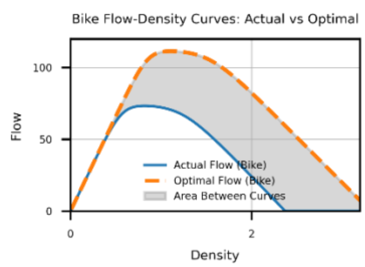

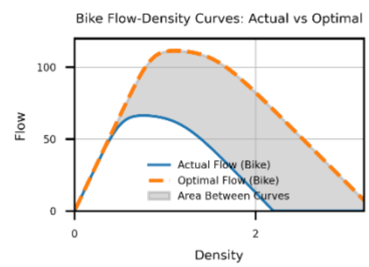

Figure 2.10 illustrate the fitted NFD for each fairness objective alongside the optimal reference to assess the ability of each objective to reach optimal states.

Existing

Utilitarian

Sufficiency

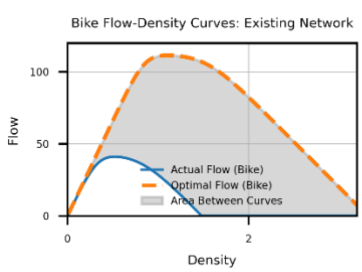

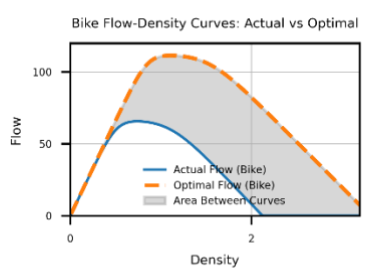

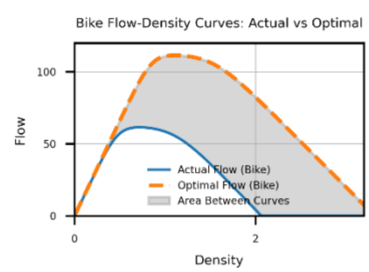

Accessibility Gap Mode

Accessibility Gap Income

Accessibility Gap Access

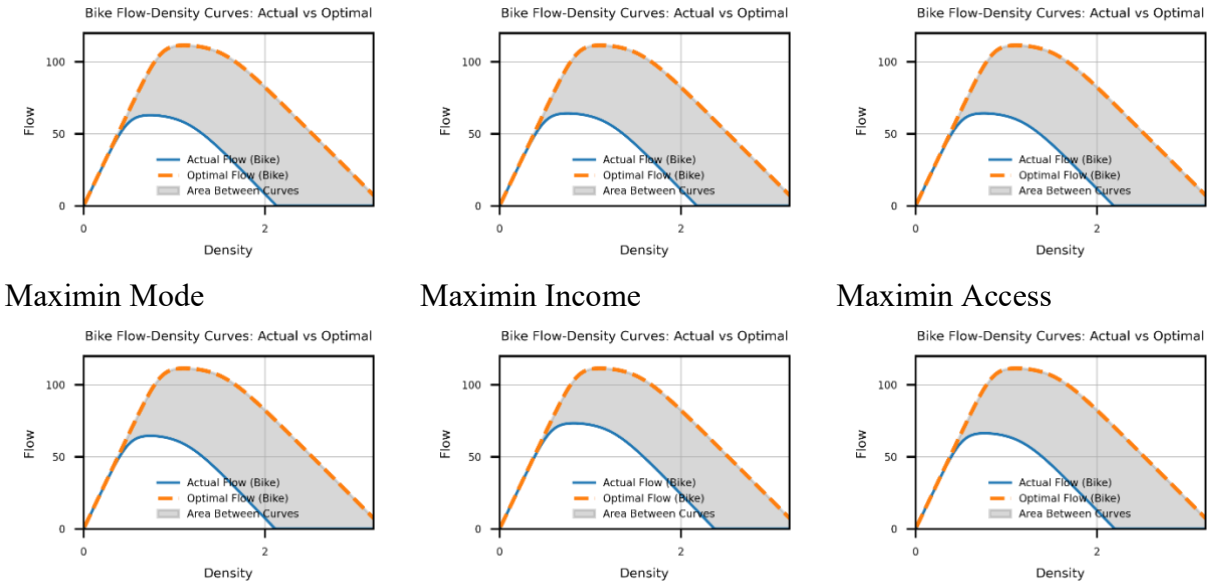

Figure 2.10 Difference Between Optimal Micromobility NFD and Best Solution from G

Across all networks, a consistent pattern is shown: the optimal flow curve (dashed orange) lies above the actual flow curve (solid blue), indicating that these solutions do not reach the optimal maximized flow. This may reflect the infeasibility of the macroscopic solution applied to the link-choice optimization. Although constraints are imposed to ensure the overall area is constant on the aggregate network-wise estimation, the optimal areas may not be able to be created through the summation of individual links.

Furthermore, the resulting GA optimized curves are strikingly similar to each other despite their alternative objectives and constraints. This pattern, as well as the relationship between budget and the optimal NFD, shows the limitations of achieving either optimality in terms of NFD deviation or differentiation across the varied ethical objectives. To complement the micromobility analysis,

Existing

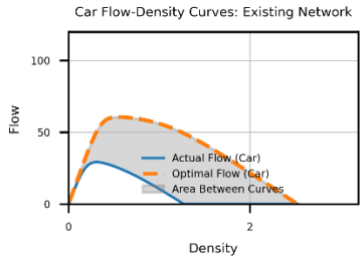

Utilitarian

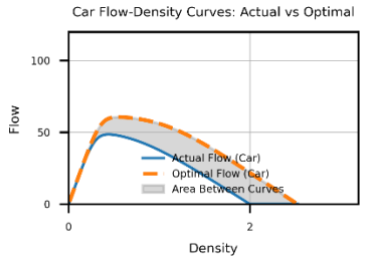

Sufficiency

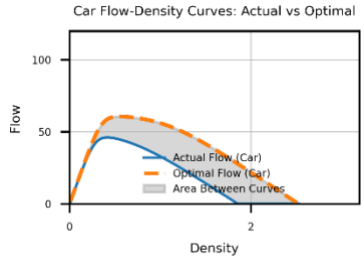

Accessibility Gap Mode

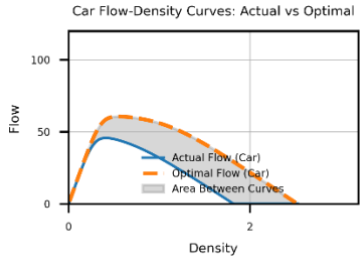

Accessibility Gap Income

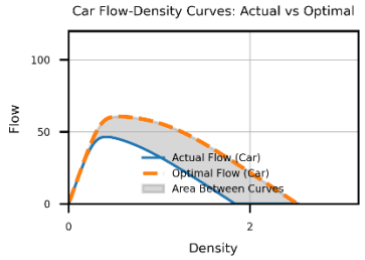

Accessibility Gap Access

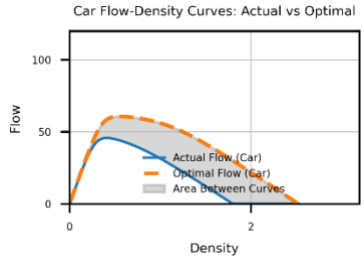

Maximin Mode

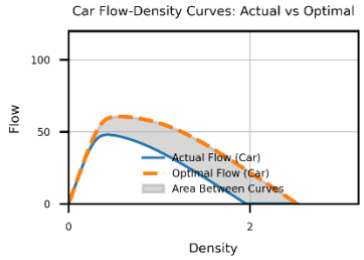

Maximin Income

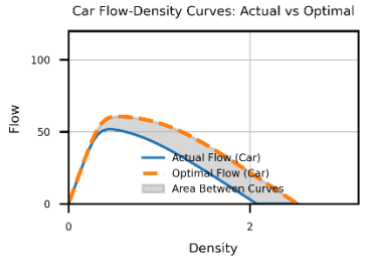

Maximin Access

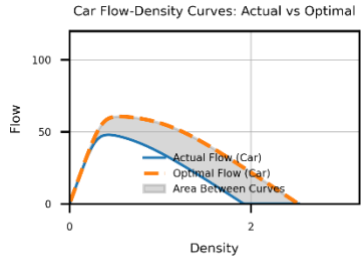

Figure 2.11 presents the fitted NFD curves for cars under each fairness objective, again compared against a reference optimal curve.

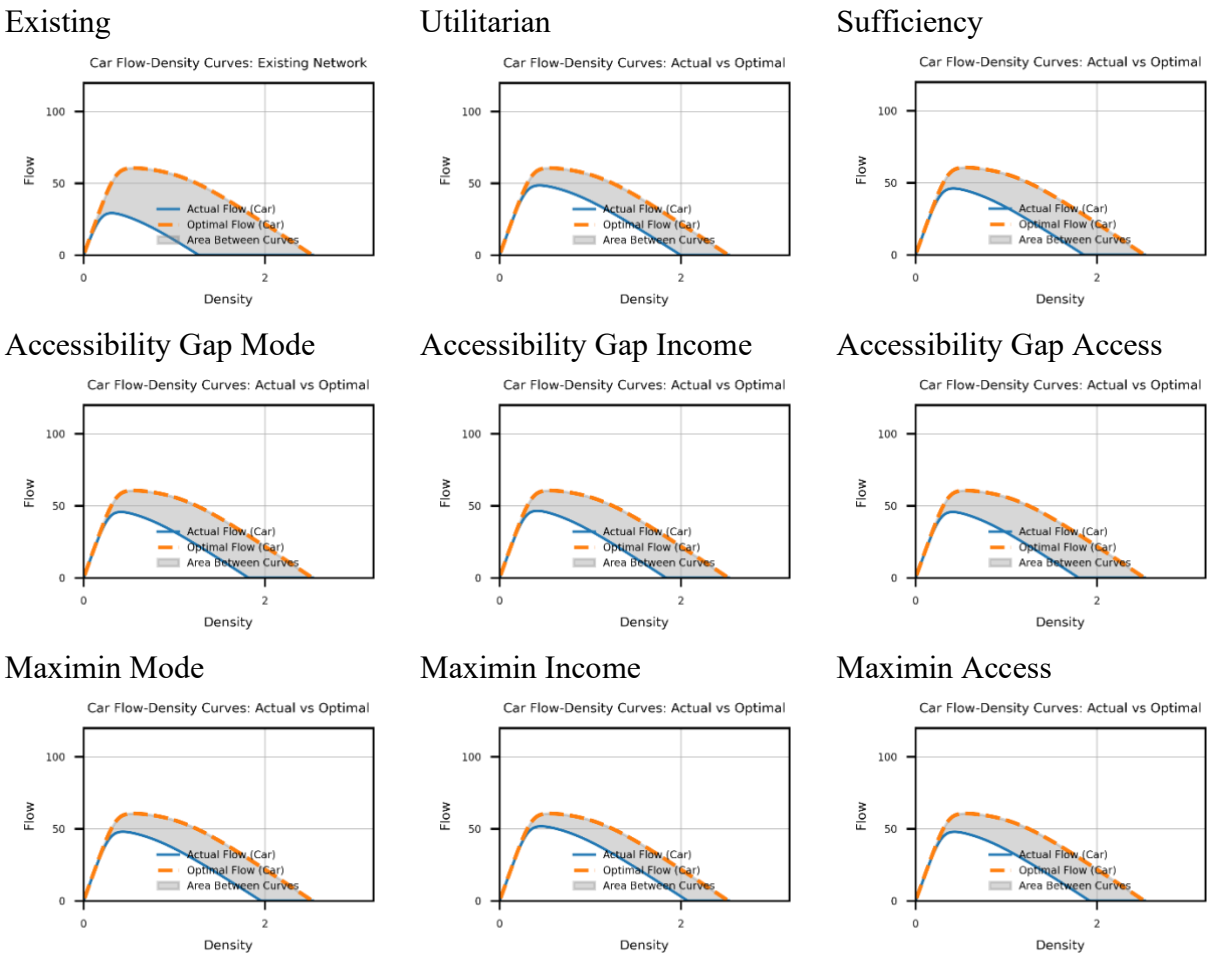

Figure 2.11 Difference Between Optimal Car NFD and Best Solution from GA

These curves follow similar patterns to micromobility vehicles, again indicating a structural incapability of the network to reach optimal due to limited physical space or budget. Compared to micromobility, the car flow curves exhibit less deviation from the optimal. Here, the maximin objectives perform well with *Maximin Mode* and *Maximin Income,* achieving the smallest area between curves, followed by the *Utilitarian* objective.

The results of each formulation in terms of the various metrics optimized for each objective are shown in Table 2.5.

Table 2.5 Evaluation Metrics of GA Results by Ethical Framework

Network	Existing		Utilitarian		Sufficiency	
	Value	Target Attainment	Value	Target Attainment	Value	Target Attainment
20-minute Access (Car)	2109	-36.21%	2109	-36.21%	2105	-36.33%
20-minute Access (Bike)	647	-80.43%	840	-74.59%	859	-74.02%
20-minute Access (Total)	2756	-58.32%	2949	-55.40%	2964	-55.17%
45-minute Access (Car)	3306	0.00%	3306	0.00%	3306	0.00%
45-minute Access (Bike)	2434	-26.38%	2714	-17.91%	2744	-17.00%
45-minute Access (Total)	5740	-13.19%	6020	-8.95%	6050	-8.50%
Low Income Access	46	-59.65%	49	-57.02%	46	-59.65%
High Income Access	37	-67.54%	40	-64.91%	39	-65.79%
Most Accessible	69	-39.47%	74	-35.09%	76	-33.33%
Least Accessible	13	-88.60%	14	-87.72%	14	-87.72%
Construction Cost ($)	-	10,000,000	9,974,336	25,664	9,791,746	208,254
Area to Optimal (Car)	72.82	-72.82	45.61	-45.61	45.59	-45.59
Area to Optimal (Bike)	182.08	-182.08	133.18	-133.18	140.54	-140.54
	Accessibility Gap (Income)		Accessibility Gap (Mode)		Accessibility Gap (Access)	
	Value	Target Attainment	Value	Target Attainment	Value	Target Attainment
20-minute Access (Car)	2109	-36.21%	2106	-36.30%	2109	-36.21%
20-minute Access (Bike)	863	-73.90%	883	-73.29%	827	-74.98%
20-minute Access (Total)	2972	-55.05%	2989	-54.79%	2936	-55.60%
45-minute Access (Car)	3306	0.00%	3306	0.00%	3306	0.00%
45-minute Access (Bike)	2718	-17.79%	2688	-18.69%	2665	-19.39%
45-minute Access (Total)	6024	-8.89%	5994	-9.35%	5971	-9.69%
Low Income Access	50	-56.14%	49	-57.02%	46	-59.65%
High Income Access	39	-65.79%	40	-64.91%	40	-64.91%
Most Accessible	74	-35.09%	74	-35.09%	71	-37.72%
Least Accessible	14	-87.72%	14	-87.72%	14	-87.72%
Construction Cost ($)	9,966,250	33,750	9,992,368	7,632	9,844,529	155,471
Area to Optimal (Car)	44.97	-44.97	45.94	-45.94	46.65	-46.65
Area to Optimal (Bike)	131.86	-131.86	135.54	-135.54	131.30	-131.30

	Maximin (Income)		Maximin (Mode)		Maximin (Access)	
	Value	Target Attainment	Value	Target Attainment	Value	Target Attainment
20-minute Access (Car)	2109	-36.21%	2109	-36.21%	2110	-36.18%
20-minute Access (Bike)	870	-73.68%	853	-74.20%	850	-74.29%
20-minute Access (Total)	2979	-54.95%	2962	-55.20%	2960	-55.23%
45-minute Access (Car)	3306	0.00%	3306	0.00%	3306	0.00%
45-minute Access (Bike)	2758	-16.58%	2685	-18.78%	2702	-18.27%
45-minute Access (Total)	6064	-8.29%	5991	-9.39%	6008	-9.13%
Low Income Access	50	-56.14%	50	-56.14%	50	-56.14%
High Income Access	40	-64.91%	40	-64.91%	36	-68.42%
Most Accessible	74	-35.09%	73	-35.96%	73	-35.96%
Least Accessible	14	-87.72%	14	-87.72%	14	-87.72%
Construction Cost ($)	11,910,409	-1,910,409	9,928,178	71,822	9,985,509	14,491
Area to Optimal (Car)	40.52	-40.52	45.87	-45.87	46.30	-46.30
Area to Optimal (Bike)	108.61	-108.61	133.41	-133.41	127.40	-127.40

All optimized networks preserve high levels of car accessibility within the 20-minute window and full accessibility within the 45-minute window, where cars can reach 3,306 accessible zones, which is the maximum accessibility. While bike access increases substantially over the existing network in all formulations, full accessibility in either window is not achieved. This increased accessibility is most significant in the *Accessibility Gap Mode* network, allowing 236 additional zones compared to the existing network to be accessible to bikes within 20 minutes. Metrics for equity improve modestly in comparison. Access for low-income zones increases in all scenarios, but only marginally. Construction costs vary by formulation, but all formulations are just below the $10M budget except *Maximin Income*, though this resulting network is also the most optimal in terms of closeness to the optimal curve for cars and bikes.

To visually compare the performance of network configurations optimized under different ethical principles, radar plots are used. Each axis represents a key evaluation metric, including total and mode-specific accessibility, construction cost, equity measures, and closeness to achieving the optimal NFD. For interpretability, all axes are scaled such that higher values indicate better performance, specifically, construction cost and NFD deviation have been inverted so that lower raw values translate into larger positions on the plot. The resulting plots are shown in Figure 2.12.

Figure 2.12 Visual Metrics of GA Solutions by Ethical Formulation

The existing network serves as a baseline, characterized by high NFD deviation and strong car accessibility, but limited bike accessibility and poor performance on equity-related metrics. The *Utilitarian* solution achieves high overall accessibility across both 20- and 45-minute thresholds. While *Sufficiency* yields the highest total 45-minute accessibility, it performs poorly in 20-minute car accessibility and low-income access. Each of the accessibility gap strategies performs best in the metric it targets, like bike accessibility, low-income access, or coverage of the least accessible zones. Among all alternatives, the *Maximin Income* solution stands out with high performance across nearly all metrics. However, this comes at the expense of a higher construction cost, reflecting a trade-off between broad improvement and budget efficiency.

To better understand the spatial implications of each fairness objective, the resulting network configurations are mapped across the grid. These maps in Figure 2.13 display the assigned lane types on each edge, representing how the GA allocates space for cars and micromobility under each formulation.

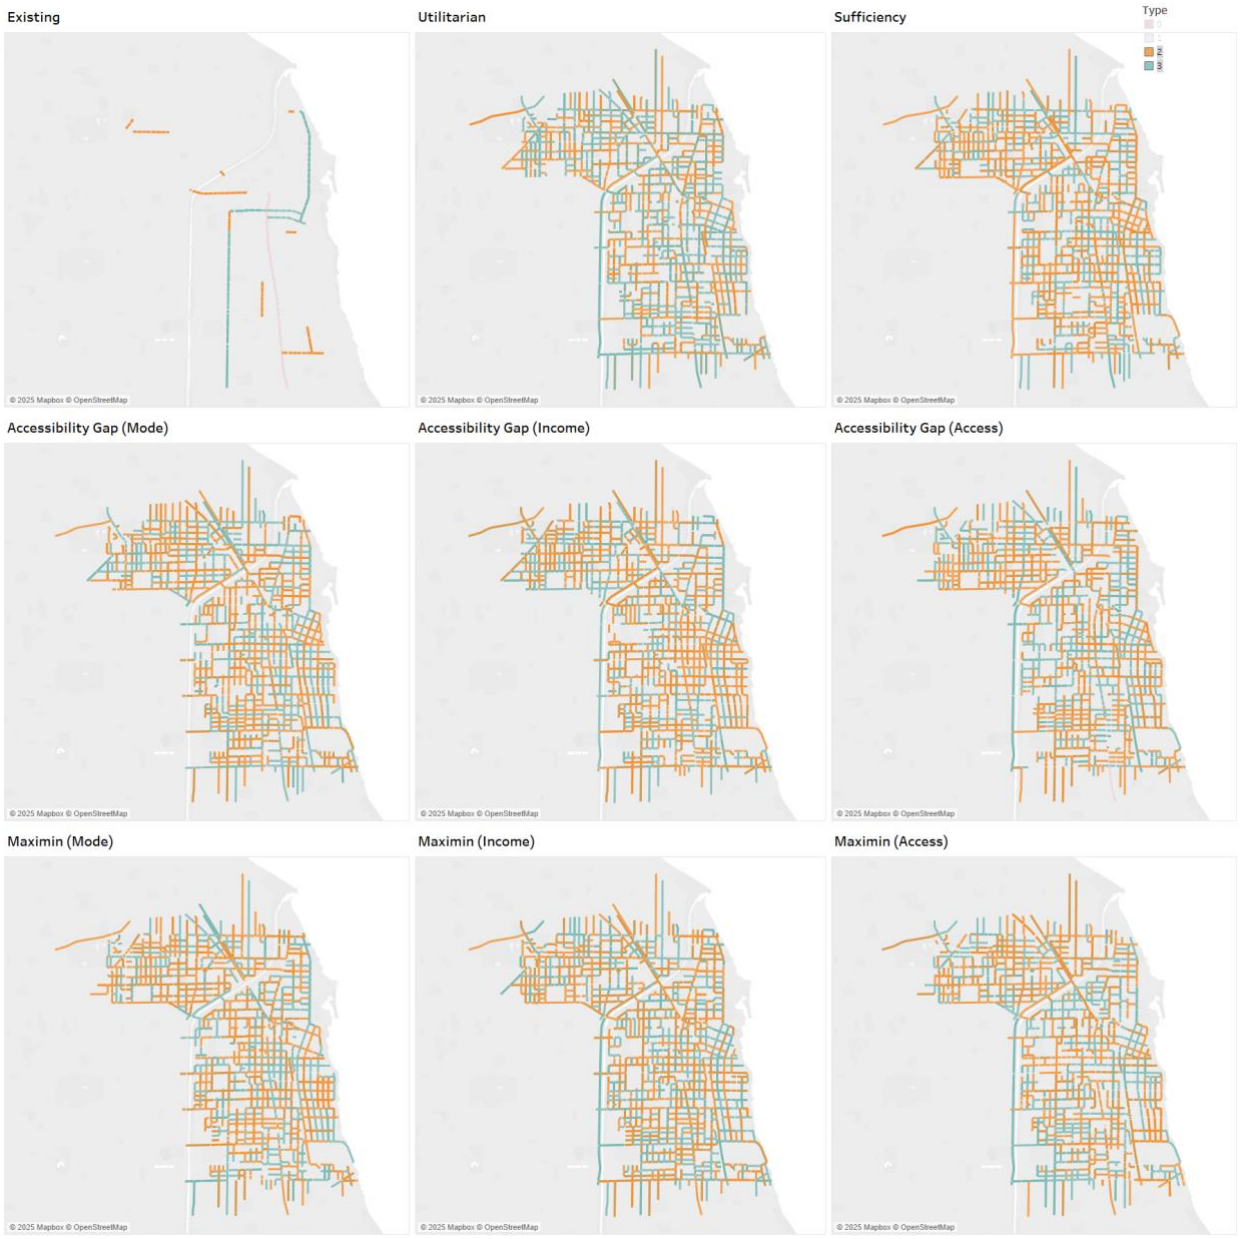

Figure 2.13 Network Representation of GA Solutions

In the existing configuration, micromobility infrastructure is sparse and fragmented, with few corridors of high-quality separation (Type 3). The existing network also includes a corridor where bikes are not allowed, identified by Type 0. In contrast, all optimized networks show expanded coverage of Types 2 and 3. *Utilitarian*, *Accessibility Gap Income*, *Accessibility Gap Access*, *Maximin Income*, and *Maximin Mode* all include corridors of bike infrastructure

connecting the western side of the network from north to south. In many of the networks, Type 3 links are concentrated in the center of the network, while Type 2 links fill in the more residential areas.

Since the objective function is designed to maximize zonal connectivity, mapping zonal access metrics allows for analysis of the spatial distribution of accessibility. These maps, shown in Figure 2.14, illustrate how accessibility varies across the network under different configurations

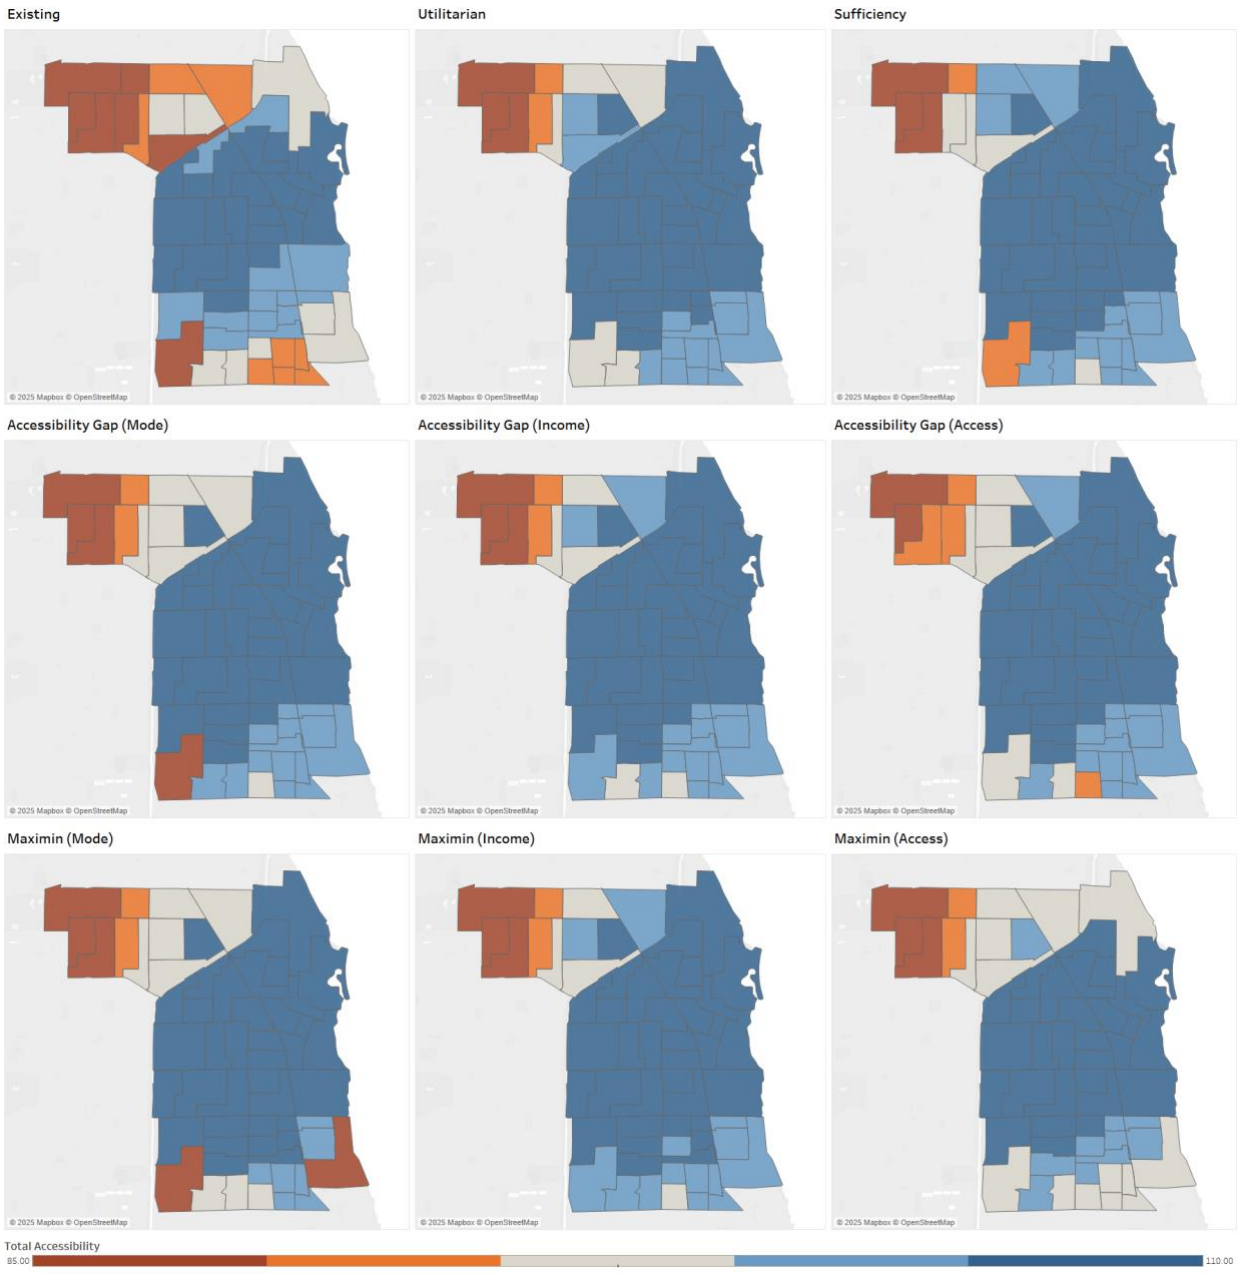

Figure 2.14 Zonal Access of GA Solutions across Fairness Objectives

The existing network shows distinct disparities between geographies, with many northern and peripheral zones in orange or beige, reflecting low access. The existing network shows high levels of connectivity in the centermost part of the map, where the downtown and most bike infrastructure is concentrated. In contrast, all optimized networks show a substantial expansion of high-accessibility areas, with more zones shaded in blue, particularly in the southern and central portions of the grid.

Utilitarian and *Sufficiency* solutions increase overall accessibility with expanded accessibility in the southern zones of the network and improvements in accessibility for the north-west portion. The modal solutions, *Accessibility Gap Mode* and *Maximin Mode* have the least uniform pattern of accessibility and are more similar to the *Existing* network but with expanded coverage in the center. *Maximin Income* and *Maximin Access* emerge as the best distribution of accessibility with the least low coverage zones of any solution.

The maps in Figure 2.15 visualize the spatial differences between each ethically guided network design and the utilitarian benchmark. To highlight only the meaningful variations, links that share the same treatment in both the displayed formulation and the *Utilitarian* solution are omitted. The full grid of both the *Utilitarian* solution and the *Existing* network are also included as reference points.

Figure 2.15 Grid Networks by Ethical Formulation Compared to Utilitarian Results

Across all formulations, there are significant variations between solutions compared to the utilitarian solution. In the *Sufficiency* scenario, a large number of links differ from the *Utilitarian* configuration, especially with Type 2 links which are used to affordably expand access, reflecting the *Sufficiency* objective's aim to raise all users to a minimum acceptable level of access. The *Accessibility Gap* strategies show more targeted variations. Both *Accessibility Gap Mode* and *Accessibility Gap Income* have some of the longest connected corridors, creating more low-stress accessibility for cyclists and connecting low-income zones. *Accessibility Gap Access* appears to prioritize more widespread coverage of the network, including more investment on the periphery. The *Maximin* principle designs demonstrate the most widespread changes relative to the *Utilitarian* baseline. *Maximin Mode* and *Maximin Income* both result in dense reallocation of

links, particularly shifting treatments in areas underserved by the *Utilitarian* plan, like the north-west portion of the map. *Maximin Access* shows the most spatially extensive reconfiguration, aligning with its goal to prioritize the zones with the lowest existing accessibility.

To assess the distributional impacts of each network configuration, cumulative distribution curves illustrate the change in 20-minute bike accessibility relative to the existing network. These plots represent the net gain or loss in access for each zone to observe which areas benefit from the reallocation of space and which ones experience reduced connectivity. By incorporating population size and median income, the curves highlight the tradeoffs between efficiency and fairness, capturing both the "winners" and "losers" under each optimization objective shown by Figure 2.16.

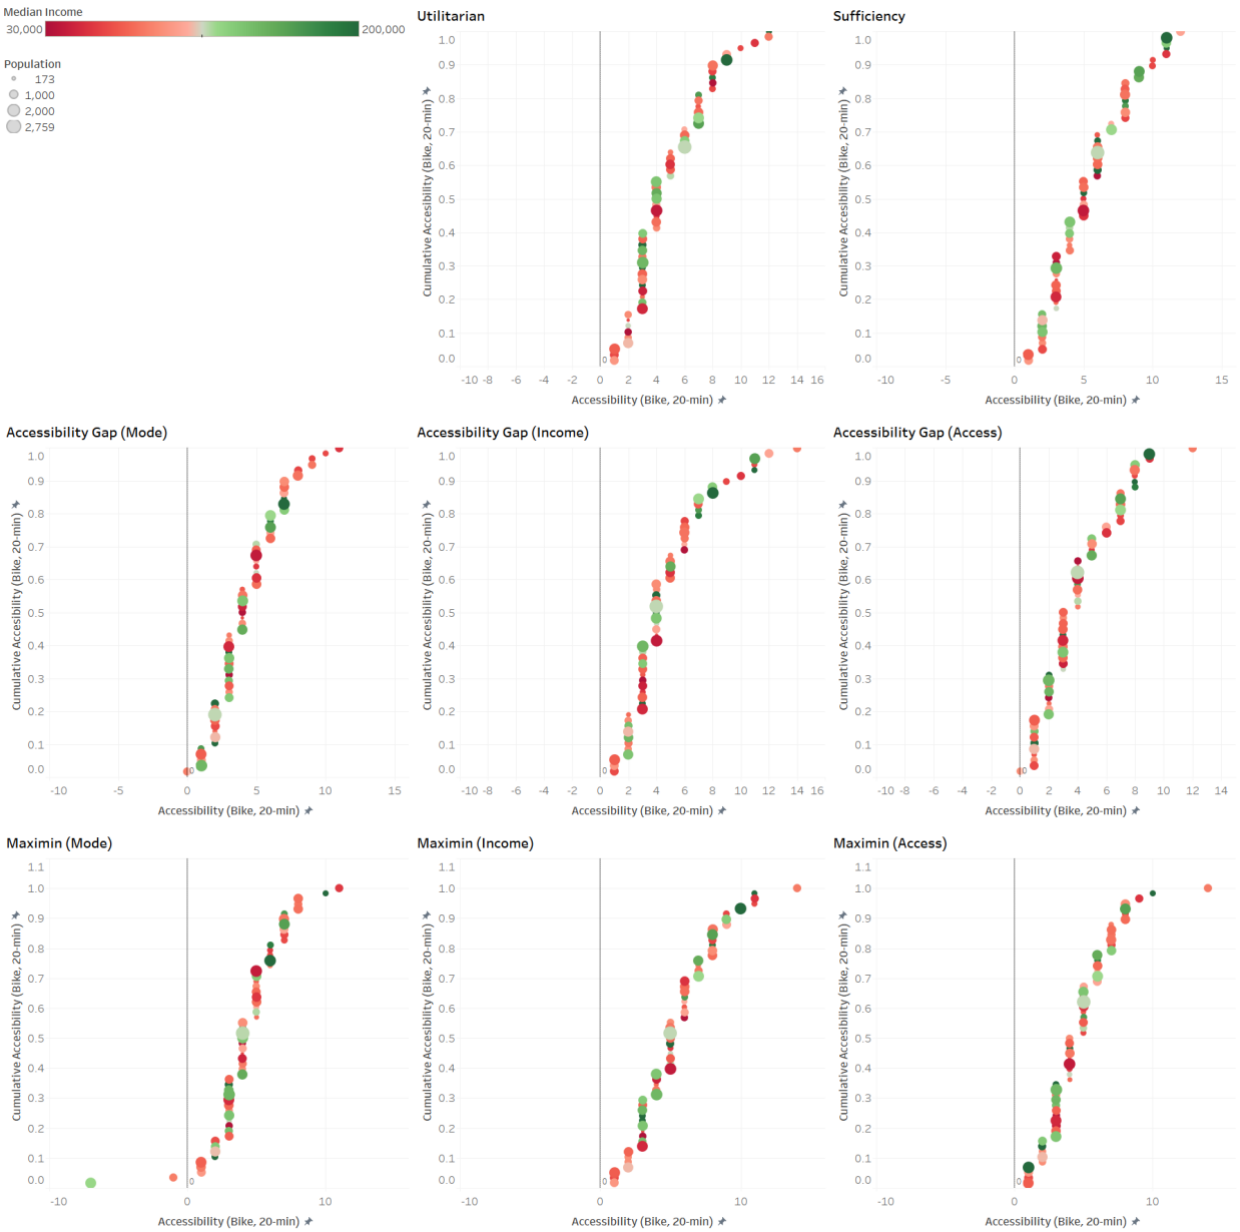

Figure 2.16 Cumulative Distribution of Improvements over Existing Infrastructure

Across all formulations, most zones experience improvements compared to the existing network. However, some tradeoffs are visible in *Maximin Mode*, where some zones lose access to benefit others. Across the distributions, there are no clear patterns in the variation in income, both low- and high-income zones see losses in different objectives, and gains are distributed to both high and low-income zones as well.

To better understand the tradeoffs between network efficiency and investment, a sensitivity analysis of the two objective function components (accessibility and flow-efficiency) over varying the construction budget was conducted. Figure 2.17 shows the resulting sensitivities of the networks toward reaching optimal flow and improvement of accessibility.

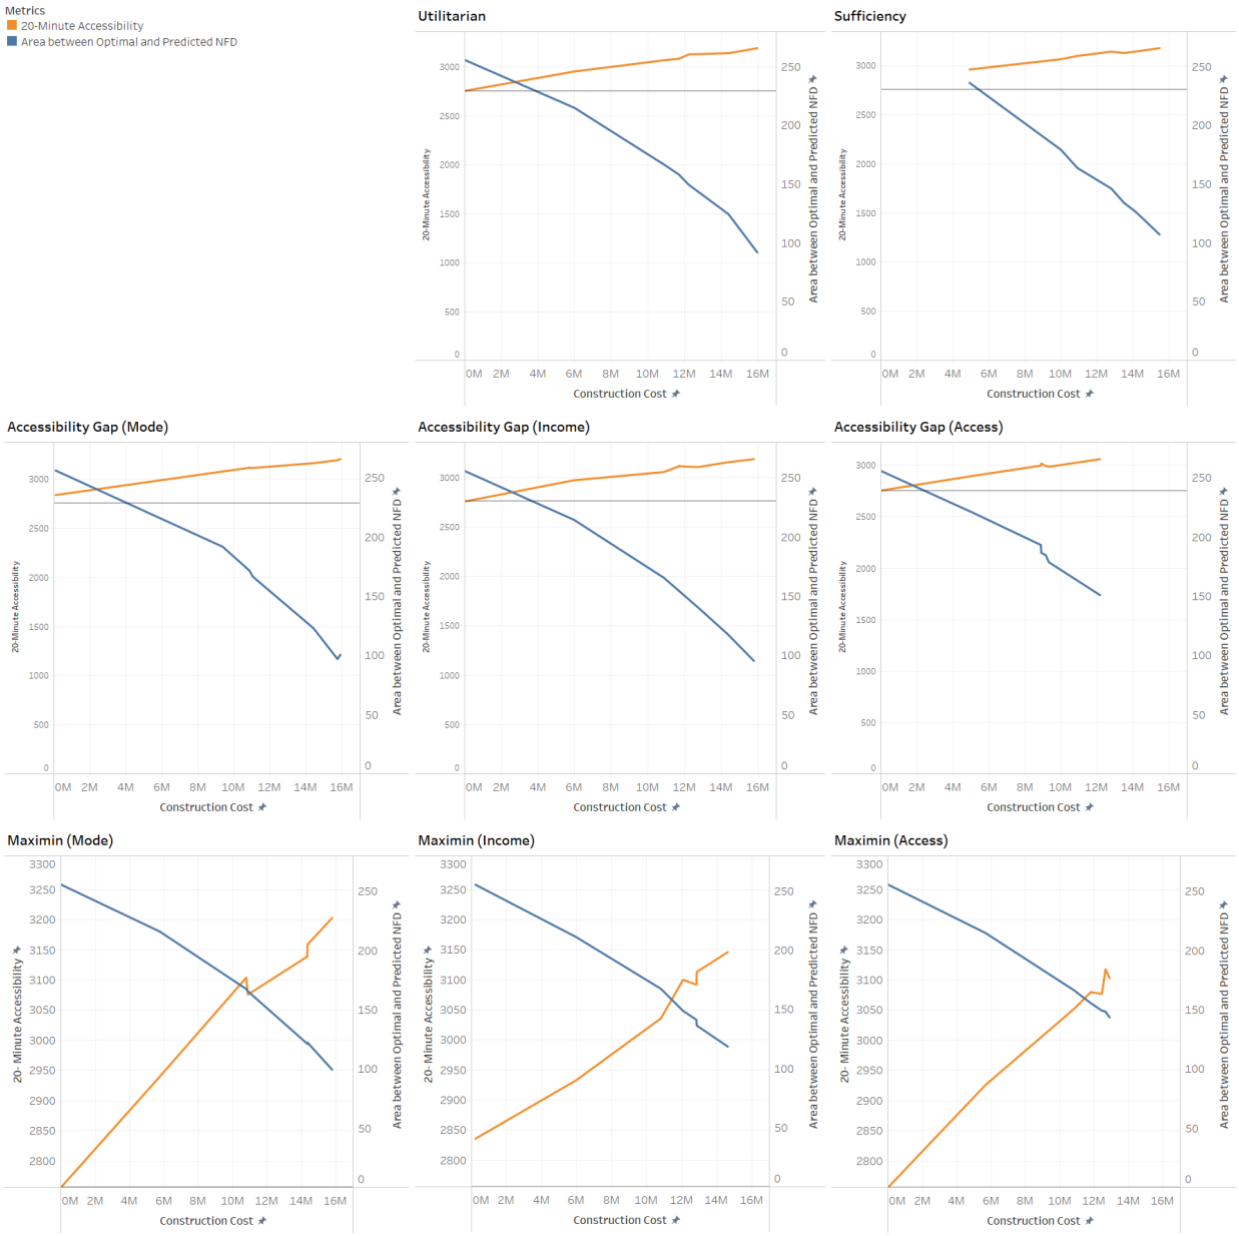

Figure 2.17 Cumulative Distribution of Improvements over Existing Infrastructure

Across all scenarios, as construction cost increases, accessibility improves and NFD deviation decreases, though no formulations achieve full accessibility or reach the optimal NFD curve. The *Utilitarian* and *Accessibility Gap Income* formulations achieve the closest fit to the optimal NFD curves. The *Accessibility Gap Mode* and *Maximin Mode* formulations achieve the highest accessibility, which is achieved through the prioritization of bike accessibility. When budgets are limited, the solution tends to preserve much of the existing network. However, even modest additions lead to noticeable improvements in accessibility compared to the baseline (grey line). For example, in *Accessibility Gap Mode*, an investment of $96,000 yields 82 more accessible zones. This improvement is achieved by adding bike lane infrastructure to the network, particularly along a major corridor that was previously exclusive to cars (Ridge Avenue).

To assess the accuracy of the GA's objective function and its underlying flow predictions, each optimized network was re-simulated. The resulting flow-density curves were compared to the predicted curves used during optimization and are shown in Figure 2.18.

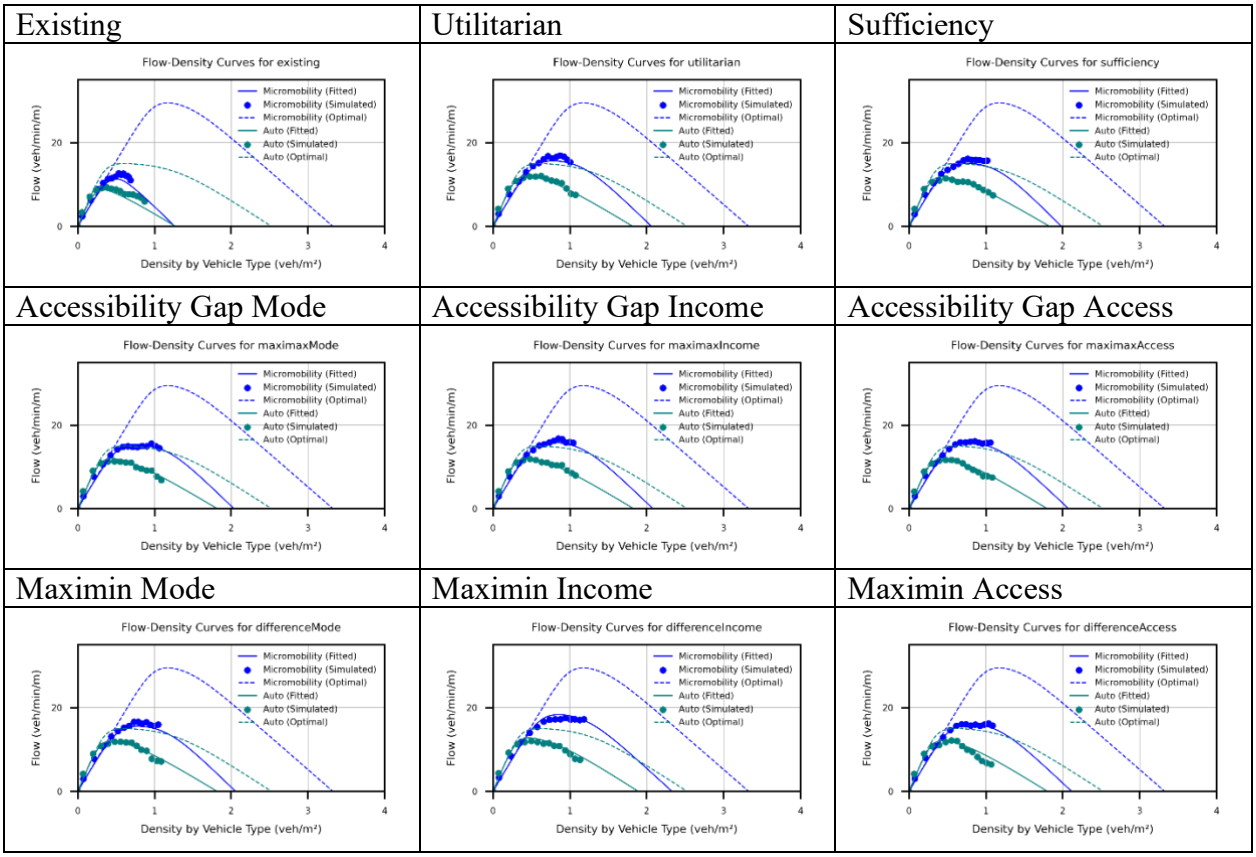

Figure 2.18 Difference Between Optimal and Actual NFD Curves and Simulated Relationships

This validation step verifies whether the simulated performance for each GA-optimized network aligns with the predicted NFD. While these graphs reinforce the finding that the GA is unable to find a network that meets the optimality defined NFD from Step 1, the predicted NFD curves and simulated points align well.

While the direct optimization network from the perspective of maximizing the function $q_m(k_m, A_m, A, I)$ is feasible in the sense of meeting the imposed constraints; these results suggest it is not feasible in the actual creation of a physical network that meets the optimal areas found, especially when enforcing budget constraints.

2.4 Summary

This chapter introduces a novel, two-step framework for the macroscopic design of bike lane networks. Building upon the network fundamental diagram (NFD) relationships identified in the prior chapter, an optimization formulation is developed that finds flow-maximizing values of the network attributes. The second step of the framework translates these optimal network characteristics onto specific links within the Evanston network using a Genetic Algorithm (GA). The GA selects configurations of links with no bike lanes, bike lanes, and separated bike lanes. The objective of the GA is to maximize accessibility while minimizing the difference between the theoretical curves of a solution network and the optimal shape found in Step 1. This solution method includes custom operators, constraint handling, and a repair function to maintain solution feasibility, improve convergence, and ensure spatial coherence of the proposed infrastructure.

To test how different ethical frameworks influence the resulting network design, the GA is implemented under multiple fairness objectives: utilitarian, sufficiency, accessibility gap, and maximin principle. Each of these formulations produces a distinct configuration of bike infrastructure with varying levels of accessibility, equity, and efficiency. Despite these differences, all optimized networks outperform the existing network, which is sparse and fragmented.

Performance metrics, convergence behavior, NFD curve comparisons, and spatial maps collectively illustrate how different objectives shape both the technical and social outcomes of the network design. Cumulative distribution curves of access gains, differentiated by income and population, reveal the extent to which each objective benefits or burdens specific zones. These analyses demonstrate that while all objectives deliver gains, tradeoffs are inevitable, and fairness criteria play a critical role in determining who benefits most.

Analysis of the resulting networks across all performance metrics highlights important tradeoffs. The *Accessibility Gap Mode* formulation yields the largest gains in 20-minute bike accessibility; however, *Maximin Income* is shown to be the best overall solution with strong results across all metrics, including accessibility, equity, and efficiency, though it slightly exceeds the allocated budget. In terms of zone-level accessibility, *Maximin Access* achieves the best outcomes, with few low-access zones in the final network. When measuring improved accessibility compared to the existing network, a small number of zones lose accessibility in the *Maximin Mode* solution.

A sensitivity analysis to construction costs provides additional insight into how different optimization goals respond to varying investment levels. Both accessibility and efficiency in terms of the area between NFD curves respond to increased budget, resulting in smaller areas between curves and higher accessibility values. Under increased budgets, the U*tilitarian* objective most closely reaches the theoretical NFD, while mode-based objectives (*Accessibility Gap Mode* and *Maximin Mode*) deliver the highest aggregate accessibility gains with additional

funding. Finally, a re-simulation of the GA-selected networks is used to validate the predictive accuracy of the functional form developed. Taken together, these results provide a set of high-performing solutions that reflect different priorities. By exploring alternative designs at a given budget, planners can present the tradeoffs between equity, efficiency, and access in transparent ways. Community input should guide the selection of objectives and final designs, ensuring that the infrastructure reflects local values and the needs of those it is meant to serve.

3 A MULTI-AGENT SIMULATION PLATFORM FOR URBAN STREETS

3.1 Introduction

Urban streets are dynamic environments where multiple types of users interact: cars, buses, bicycles, pedestrians, delivery vehicles, rideshare vehicles, etc. These interactions shape safety, efficiency, and overall street functionality. Modeling such complexity requires a framework that captures both predictable flows and stochastic events. This chapter presents a multimodal urban traffic modeling framework, designed to simulate the interplay between flowing and stopping agents and their interactions. By conceptualizing traffic as a combination of flows, stops, and localized interactions, the framework allows realistic simulation of street-level dynamics and assessment of design and operational performance.

Street dynamics operate across multiple spatial scales. At the city-wide scale, models focus on broad operations, such as rideshare fleet movement, public transit network design, and vehicle repositioning, providing insights into demand concentrations and likely flow corridors. At an intermediate scale, such as a neighborhood or corridor, sub-network models capture more localized features, including bike lane design, transit stop placement, signal timing, and shared mobility dock locations. These models reveal nuances in how different street users interact in specific areas. At the street-level scale, attention is on micro-scale interactions, such as a double-parked delivery truck, a bus pulling into a stop, pedestrians crossing mid-block, or cyclists navigating around obstacles. While higher-level models provide expected flows, the precise timing and locations of these events are stochastic, driven by variability in where and when agents stop and by the interactions among them, which may cause delays, conflicts, or rerouting.

Urban street users can be classified as either flowing or stopping agents. Flowing agents primarily travel through the street to reach their destinations, including vehicles moving along a corridor or cyclists riding along a bike lane. Stopping agents, in contrast, treat the street as a destination or task location; examples include rideshare pickups, buses pausing at stops, delivery trucks temporarily stopping, or pedestrians crossing mid-block. Interactions between flowing and stopping agents are central to the functionality and safety of streets, necessitating a modeling framework that can represent these interactions accurately.

3.2 Simulation Framework

The framework represents flowing agents with distributions of speeds and flow rates to capture natural variability and stopping agents with stochastic stop locations and dwell times. It incorporates street infrastructure such as lane widths, sidewalks, bike lanes, and curb allocations, which influence agent movement and interactions and affect throughput, safety, and overall user experience. The framework simulates complex interactions, including lane blockage by delivery trucks, unexpected pedestrian crossings, or bus stop delays, allowing evaluation of congestion, safety, and efficiency. By integrating stochastic demand, infrastructure features, and interaction rules, this simulator provides a tool for understanding and testing urban street designs under realistic conditions.

Each agent in the simulation is represented as an independent entity with its own decision-making capabilities. Agent types include autonomous vehicles (AVs), connected vehicles (CVs), human-driven vehicles (HDVs), bicycles, and pedestrians. Roadways can be configured flexibly; for a 60 ft right-of-way, approximately 200 different layouts are possible, specifying vehicle lane widths and directions, bike lane placement and width, and pedestrian sidewalk dimensions. The

simulator models hierarchical agent movement, including longitudinal and lateral dynamics, infrastructure-related decision processes, and multi-modal interactions. Figure 3.1 outlines some of these layouts.

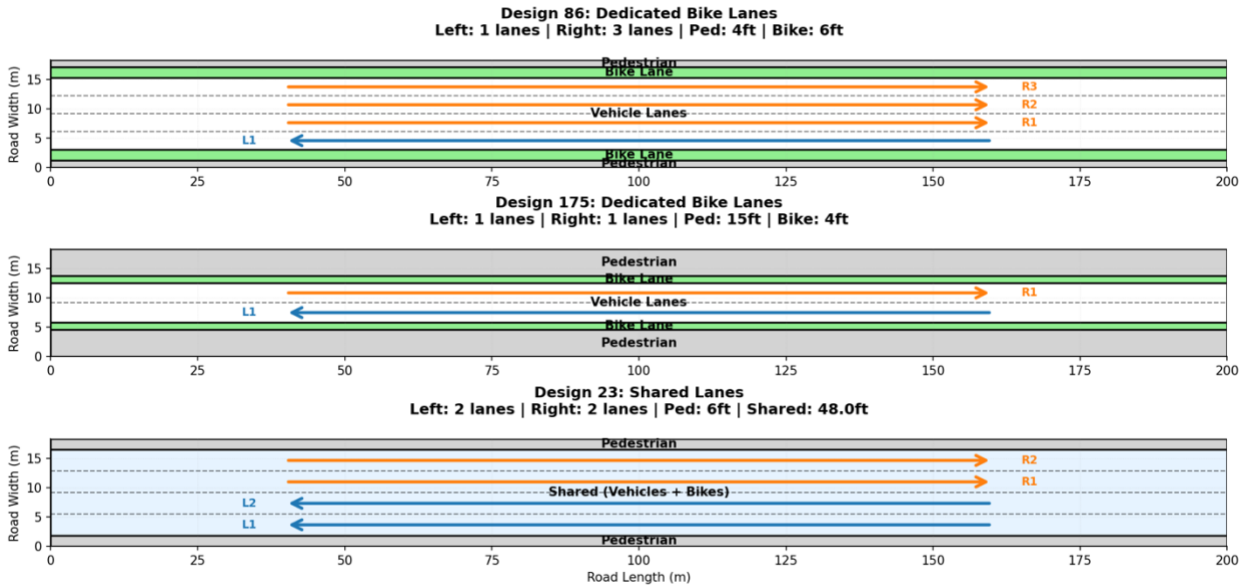

Figure 3.1 Various Simulation Street Layouts

3.3 Base Movement Models

A number of models are implemented in this simulator to represent agent movement. Some of them are borrowed from the literature, others are novel models. Figure 3.2 show these models implemented from the literature.

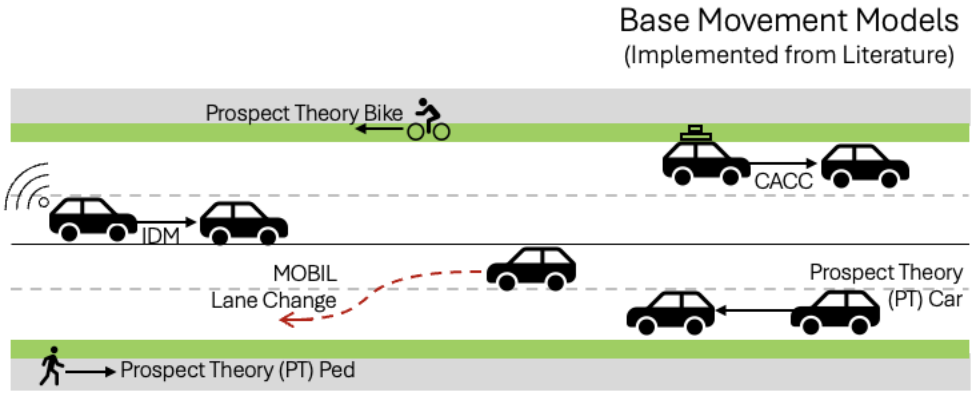

Figure 3.2 Movement Models from Literature

3.3.1 Automated Vehicle Control Model

Automated vehicles are modeled using a longitudinal control framework inspired by [4] where vehicles continuously sense the environment and react to leader motion. The effective stopping distance for vehicle $i$following $i-1$is:

$$\Delta x_n = (x_{i-1} - x_i - l_{i-1}) + v_i\tau + \frac{v_{i-1}^2}{2a_{i-1}^{decc}} \quad \text{(Eq 3.1)}$$

where x_i is the longitudinal position of vehicle $i$in meters, $v_i$is its speed in meters per second, l_{i-1} is the length of the leading vehicle, $\tau$is the driver response delay in seconds, and a_{i-1}^{dec} is the maximum comfortable deceleration of the leader in meters per second squared. Because the sensing system has limited range, the usable stopping distance is

$$\Delta x = min(SensorRange, \Delta x_n) \quad \text{(Eq 3.2)}$$

The maximum safe speed that ensures stopping within this distance is given by:

$$v_{max} = \sqrt{-2a_i^{decc}\Delta x} \quad \text{(Eq 3.3)}$$

where $a_i^{\text{dec}}$is the maximum deceleration of vehicle i. The vehicle's desired acceleration incorporates the leader's acceleration a_{i-1}, the relative velocity $(v_{i-1}-v_i)$, and the difference between the current spacing $s_i$and the desired spacing s_{ref}:

$$a_i^d(t) = k_a\,[a_{i-1}(t-\tau)] + k_v[v_{i-1}(t-\tau) - v_i(t-\tau)] + k_d[s_i(t-\tau) - s_{ref}] \quad \text{(Eq 3.4)}$$

Here, k_a, k_v, and $k_d$are control gains, and $s_{\text{ref}}$is the desired spacing, calculated as the minimum of three components: the minimum gap s_{min}, the system-based spacing $s_{\text{system}} = v_i\tau$, and the safe-stopping spacing $s_{\text{safe}} = v_{i-1}^2/[2(1/a_i^{\text{dec}} - 1/a_{i-1}^{\text{dec}})]$. The actual acceleration applied is bounded to ensure compliance with the safe-speed limit:

$$a_i(t) = min\left(a_i^d(t), k\big(v_{max} - v_i(t)\big)\right) \quad \text{(Eq 3.5)}$$

where $k$moderates convergence toward the safe-speed boundary. This formulation guarantees that AVs maintain safe headways while responding smoothly to surrounding traffic.

3.3.2 IDM for Connected Vehicles

Connected Vehicle and longitudinal dynamics follow the Intelligent Driver Model (IDM) [5]:

$$a = a_0\left(1 - \left(\frac{v}{v_0}\right)^{\delta} - \left(\frac{s^*(v,\Delta v)}{s}\right)^2\right) \quad \text{(Eq 3.6)}$$

where $a_0$is the maximum acceleration, $v$the current speed, $v_0$the desired free-flow speed, $s$the gap to the leader, and δan acceleration exponent. The desired dynamic spacing is:

$$s^*(v,\Delta v) = s_0 + vT + \frac{v\Delta v}{2\sqrt{a_0 b}} \quad \text{(Eq 3.7)}$$

with $s_0$the minimum gap, $T$the desired time headway, $\Delta v = v - v_{\text{leader}}$ the relative speed, and $b$the comfortable deceleration.

3.3.3 Multi-Regime Prospect-Theory based Car-Following Behavior

Human drivers evaluate candidate accelerations using Prospect Theory (PT) [6]. Two behavioral regimes are distinguished: uncongested (UC) and congested (C) [7]. The expected utility of a candidate acceleration $a_n$is

$$U_{PT}(a_n) = P(C)\cdot U_{PT}^{C}(a_n) + P(UC)\cdot U_{PT}^{UC}(a_n) \quad \text{(Eq 3.8)}$$

where $U_{PT}^{C}$and $U_{PT}^{UC}$are regime-specific utilities, and $P(C), P(UC)$are probabilities derived from local traffic density and spacing. Collisions are accounted for as

$$U_{PT}(a_n) = \left(1 - p_{n,i}\right) U_{PT}(a_n) - p_{n,i}\, k(v,\Delta v) \quad \text{(Eq 3.9)}$$

where $p_{n,i}$is the probability of collision and $k(v,\Delta v)$is the crash severity function. Final accelerations are stochastically sampled from

$$f(a_n) \propto \exp\left(\beta_{PT}\, U_T(a_n)\right) \quad \text{(Eq 3.10)}$$

3.3.4 Pedestrian and Bike Prospect Theory Model

Pedestrians and bicycles select speed $v$and heading $\omega$to maximize expected utility while avoiding collisions [8]. Candidate velocities are evaluated using the subjective value function

$$SV(v,\omega) = (\eta \cdot \text{align})^{sp\_ratio^{\xi}} \quad \text{(Eq 3.11)}$$

where $\text{align} = \max(0, \cos(\omega - \text{goal}_{\omega}))$, $sp\_ratio = v/v_{\text{pref}}$, $\eta$is a utility weight, and ξcontrols sensitivity to deviations from the preferred speed v_{pref}. Expected utility accounts for collision probability $p_{\text{col}}$and collision weight W_c:

$$U_{\text{expected}} = (1 - p_{\text{col}}) \cdot SV - p_{\text{col}} \cdot W_c \quad \text{(Eq 3.12)}$$

Velocities are smoothed with a first-order lag

$$v_{\text{new}} = v + \frac{\Delta t}{\tau} (v_{\text{desired}} - v) \quad \text{(Eq 3.13)}$$

and stochastic noise is added:

$$v_{\text{final}} = v_{\text{new}} + \mathcal{N}(0, \sigma\sqrt{\Delta t}) \quad \text{(Eq 3.14)}$$

where $\tau$is the smoothing time constant and σrepresents process noise.

3.3.5 MOBIL Lane Changing Model

Lane changes follow the MOBIL model [9]. A lane change is allowed if it satisfies both a safety and incentive criterion. Safety requires that the follower in the target lane does not exceed a safe braking limit b_{safe}:

$$a'_{B'} > b_{safe} \quad \text{(Eq 3.15)}$$

where $a'_{B'}$is the acceleration of the back vehicle after the change. The incentive criterion requires a net advantage to the driver:

$$a'(M') - a(M) > p\,[a(B') - a'(B')] + a_{\text{thr}} \quad \text{(Eq 3.16)}$$

where $a_M$and $a'_{M'}$are the ego vehicle accelerations before and after, $p$is the politeness factor, and $a_{\text{thr}}$is a minimum advantage threshold.

3.4 Novel Model Implementations

A number of models are newly implemented in this simulator to represent agent movement. show these novel models.

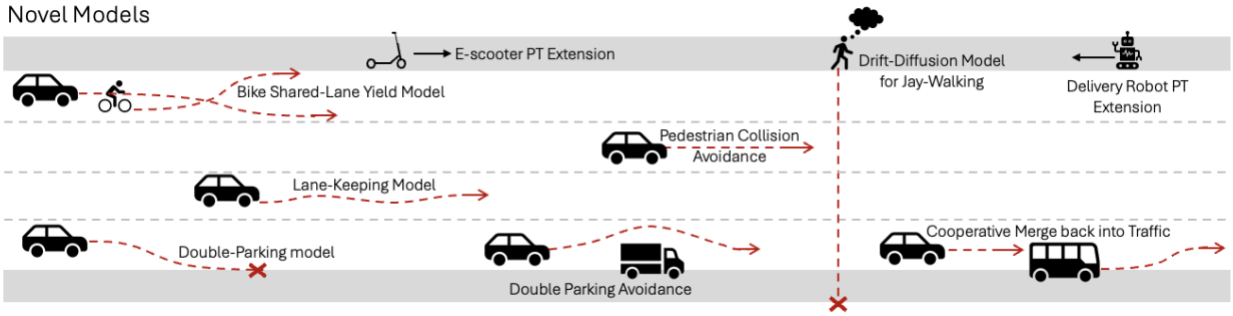

Figure 3.3 Novel Movement Models

3.4.1 Lane Keeping Dynamics Model

Accurately modeling how vehicles stay within their lanes is essential for representing traffic flow in urban environments. While longitudinal behavior such as acceleration and car-following has been widely studied, lane-keeping behavior, the small, continuous lateral adjustments drivers make to remain centered in a lane, requires separate consideration. These adjustments are influenced not only by the physical boundaries of the lane, but also by nearby vehicles and the driver's ability to control lateral motion.

This section presents a physics-based lane-keeping model that represents lateral vehicle motion as the result of three interacting effects: a restoring force that pulls the vehicle toward the lane center, repulsive forces from nearby vehicles, and damping that limits abrupt lateral movement. Together, these components produce realistic within-lane motion over time.

$$v_{y,i}(t+1) = \gamma v_{y,i}(t) + \left[F_{center,i}(t) + F_{neighbor,i}(t) + \xi_i(t)\right]\Delta t \qquad \text{(Eq 3.17)}$$

Here, $v_{y,i}(t)$denotes the lateral velocity of vehicle i, $\gamma \in [0,1]$is a damping coefficient that controls how quickly lateral motion decays, and $\Delta t$is the simulation time step. The term $\xi_i(t)$represents Gaussian noise with zero mean and variance σ^2, capturing stochastic variation in human driving behavior. The lane-centering force pulls the vehicle toward the center of its lane and is proportional to its lateral deviation:

$$F_{center,i}(t) = -\beta_{center} d_i(t) \qquad \text{(Eq 3.18)}$$

where $\beta_{\text{center}} > 0$is the centering force coefficient and $d_i(t)$is the signed lateral distance between the vehicle's current position and the lane centerline. The negative sign ensures that deviations to one side generate forces in the opposite direction, producing a stabilizing effect similar to a spring. To maintain lateral separation from nearby vehicles, the model includes a neighbor repulsion force that decreases with distance:

$$F_{neighbor,i}(t) = \beta_{neighbor} \sum_{j \in N_i(t)} sgn\left(y_i(t) - y_j(t)\right) / r_{ij}(t)^2 \qquad \text{(Eq 3.19)}$$

In this expression, $\beta_{\text{neighbor}} > 0$controls the strength of repulsion, $\mathcal{N}_i(t)$is the set of neighboring vehicles, and $r_{ij}(t)$is the Euclidean distance between vehicles $i$and j. The sign function ensures that the force pushes vehicles laterally away from one another.

Together, these equations describe a stable and flexible lane-keeping model that captures realistic lateral behavior while remaining computationally efficient for large-scale traffic simulations.

The model is calibrated using real-world trajectory data of human drivers from the Third Generation Simulation (TGSIM) dataset collected at George Washington University's Foggy Bottom campus in Washington, D.C. [10].Through multi-objective optimization using genetic algorithms, parameter sets are identified that best reproduce observed lateral velocities, achieving an acceptable fit on validation data. While the current calibration focuses on human driving behavior, the model framework is designed to accommodate autonomous vehicle (AV) behavior in future simulation studies through parameter adjustment, particularly reduced stochastic noise (σ) to reflect more precise lateral control.

3.4.2 Car/Ped Stopping Model

Vehicles approaching stops like rideshare pickup or drop-off, buses, or delivery vehicles approaching stops decelerate smoothly:

$$a = \min\left(\frac{v^2}{2d_{\text{remaining}}}, a_{\max}\right) \quad \text{(Eq 3.20)}$$

where $v$is current speed, $a_{\max}$the maximum comfortable deceleration, and $d_{\text{remaining}}$the distance to the stopping point. Lateral adjustments are applied for the car to pull-over or double park:

$$y_{\text{new}} = y + v\sin(\theta)\,\Delta t \quad \text{(Eq 3.21)}$$

with $\theta$the drift angle toward lane edge.

3.4.3 Drift-Diffusion Model (DDM) for Pedestrian/Bicycle Crossing Decisions

The drift diffusion model Implementation is as in the Phase I report for this project [11].

3.4.4 Shared-Lane Yielding Model

The Shared Lane Yield (SLY) Model simulates interactions between bicycles and cars in non-traditional shared lanes, where bicycles may laterally adjust to allow vehicles to pass. Unlike conventional lanes, lateral movement is used to resolve potential conflicts rather than strict lane discipline. In the SLY framework, bicycles detect approaching vehicles and yield toward the edge of the lane if they are putting the car at a significant disadvantage. Cars, in turn, continue forward, passing around the bicycle without slowing. Lateral yielding and passing movements are executed gradually using a small drift angle, ensuring smooth and realistic motion

Bicycles decide to yield based on the potential disadvantage imposed on a following vehicle, measured as the reduction in expected acceleration if the bicycle were present:

$$disadvantage = |a_{with\ bike} - a_{without\ bike}| \quad \text{(Eq 3.22)}$$

If this value exceeds a threshold, the bicycle politely drifts toward the lane edge, allowing cars to pass safely. The lateral motion toward the lane edge is governed by:

$$x_{new} = x + v \cdot cos(\theta)\, \Delta t\ , \quad y_{new} = y + v \cdot sin(\theta)\, \Delta t \quad \text{(Eq 3.23)}$$

where $v$is the agent's forward speed and $\theta$is a fixed drift angle (10°) toward the target lateral position. Longitudinal acceleration is updated according to the forward speed.

Bicycles move laterally to avoid impeding cars, while cars maintain their course and speed, passing on the open side of the lane. Once the vehicle has passed, the bicycle returns to the lane center smoothly. This model captures realistic bicycle–car interactions in shared lanes, where bicycles yield laterally while vehicles pass, using smooth drift dynamics to ensure safety and continuity of motion.

3.4.5 Stopped Vehicle Avoidance Model

Vehicles approaching stopped cars follow a hierarchical navigation strategy, prioritizing lane changes, then lateral avoidance if necessary, and incorporating cooperative behavior for merging vehicles. Each vehicle evaluates movement priorities in the following order: conflict resolution, stopping behavior, discretionary lane changes using MOBIL, shared lane yielding, and stopped vehicle avoidance as a fallback. When possible, vehicles attempt a discretionary lane change to pass the stopped obstacle, considering safety and acceleration benefits.

If a lane change is not feasible, lateral avoidance is applied by drifting to the side of the lane to maintain a lateral clearance from the stopped car. Vehicles adjust acceleration to maintain a passing speed. The model uses a three-state machine: *Approaching* in which the vehicle drifts toward the target lateral position; *Passing*, maintaining the lateral clearance alongside the stopped vehicle; and *Returning*, in which the vehicle returns to the lane center once past the stopped car by 5m longitudinally.

Cooperative behavior occurs when the stopped vehicle is in a *waiting-to-merge* state. Approaching vehicles reduce speed if feasible; otherwise, they navigate around the obstacle. If the slowdown creates a sufficient 3-second gap, the stopped vehicle merges back into traffic. Once past the obstacle, vehicles clear the avoidance state and resume normal car-following behavior.

Together, these models implement a hierarchical navigation system in which vehicles first attempt lane changes to pass pulled-over cars, fallback to lateral avoidance within the lane if lane changes are not possible, and apply cooperative strategies when interacting with stopped or merging vehicles. This is a unique feature of the simulator, as most traffic models assume a simple FIFO behavior where stopped vehicles form a queue behind them. By allowing vehicles to navigate around stopped agents on links, the model more realistically represents urban driving behavior and produces credible macroscopic traffic patterns.

3.5 Hierarchy of Agent Movements

Given the substantial number of potential movement models, this simulator uses a hierarchical structure for how agents choose their movements. The agent movement model implements this hierarchical decision structure at every simulation time step. For each agent, candidate actions are evaluated sequentially across the four layers, with higher-priority layers preempting lower ones. If a safety response is triggered, such as emergency braking or evasive lateral motion, that action is executed immediately and no further decisions are considered for that time step. If no safety-critical condition is present, agents with localized objectives, such as reaching a curbside stop, transit platform, or crossing location, enter the strategic layer, where route-aligned lane choices and deceleration plans are generated to reach the target safely. In the absence of strategic demands, the tactical layer evaluates discretionary behaviors, including lane changes governed by incentive and safety criteria, yielding in shared spaces, or lateral avoidance of stopped vehicles. When none of the higher layers produce an action, agents default to the operational layer, where longitudinal and lateral motion are determined by their base movement models, such as car-following, lane-keeping, or pedestrian and bicycle velocity selection. This hierarchical evaluation is performed independently for each agent at every simulation time step, ensuring that immediate safety is always prioritized while allowing longer-horizon objectives and discretionary behaviors to emerge naturally from the interaction of agents and infrastructure. This is outlined in Figure 3.4.

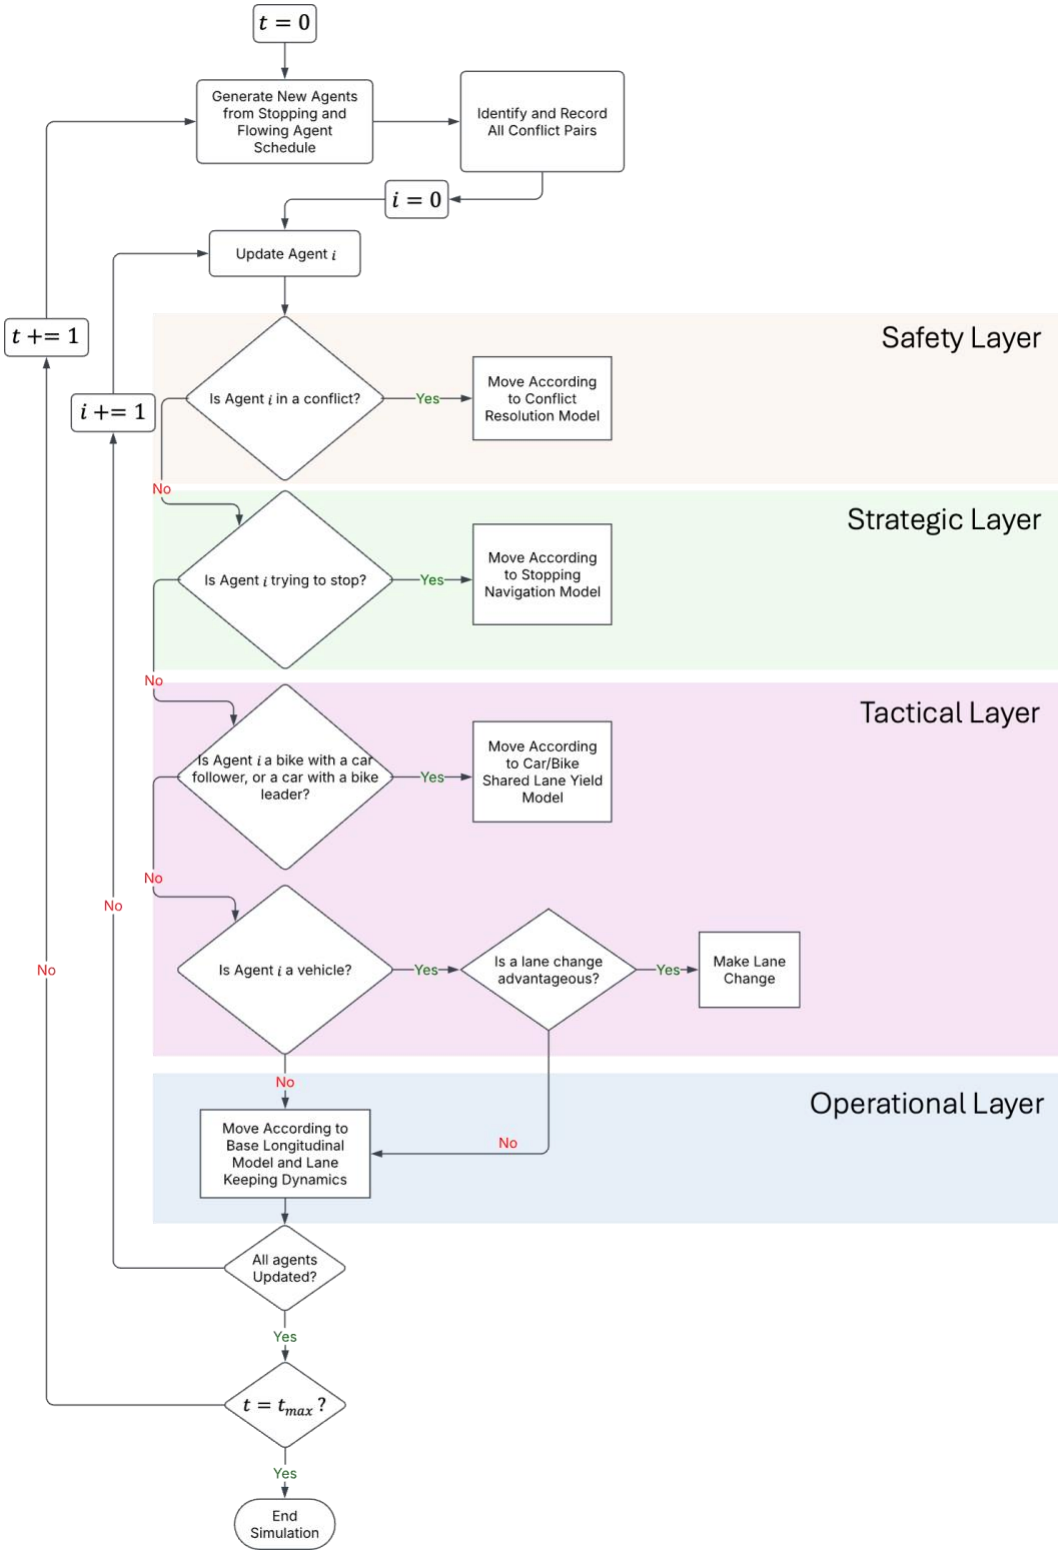

Figure 3.4 Agent Movement Model Hierarchy

3.6 Validation of Macroscopic Relationships

While the preceding sections define the microscopic motion rules for autonomous vehicles (AVs), connected vehicles (CVs), and human-driven vehicles (HDVs), the primary value of these models lies in the macroscopic traffic phenomena that emerge from their interaction. To illustrate this, a simple bidirectional corridor is simulated in which vehicles enter at a constant flow rate, and five pedestrian jaywalking events are introduced as localized disturbances. As vehicles decelerate to yield to pedestrians, these disturbances propagate upstream as backward-moving shockwaves in the space–time trajectories. Despite operating under the same geometric and demand conditions, the three vehicle classes exhibit clearly separated speed–density relationships, reflecting differences in control precision, reaction time, and headway selection. Autonomous vehicles recover most rapidly from pedestrian-induced slowdowns due to conservative gap-keeping and responsive control, connected vehicles exhibit moderate recovery as information sharing reduces but does not eliminate variability, and human-driven vehicles recover the slowest, with the greatest dispersion in speeds and headways. These differences are visible both in trajectory diagrams (Figure 3.5) and in aggregated speed–density plots (Figure 3.6), where pedestrian crossings appear as localized deviations whose dissipation rates vary by vehicle type. Overall, the results demonstrate that the proposed microscopic rules generate realistic macroscopic behavior, including shockwave formation, heterogeneous recovery dynamics, and mode-specific speed–density patterns consistent with established traffic-flow theory.

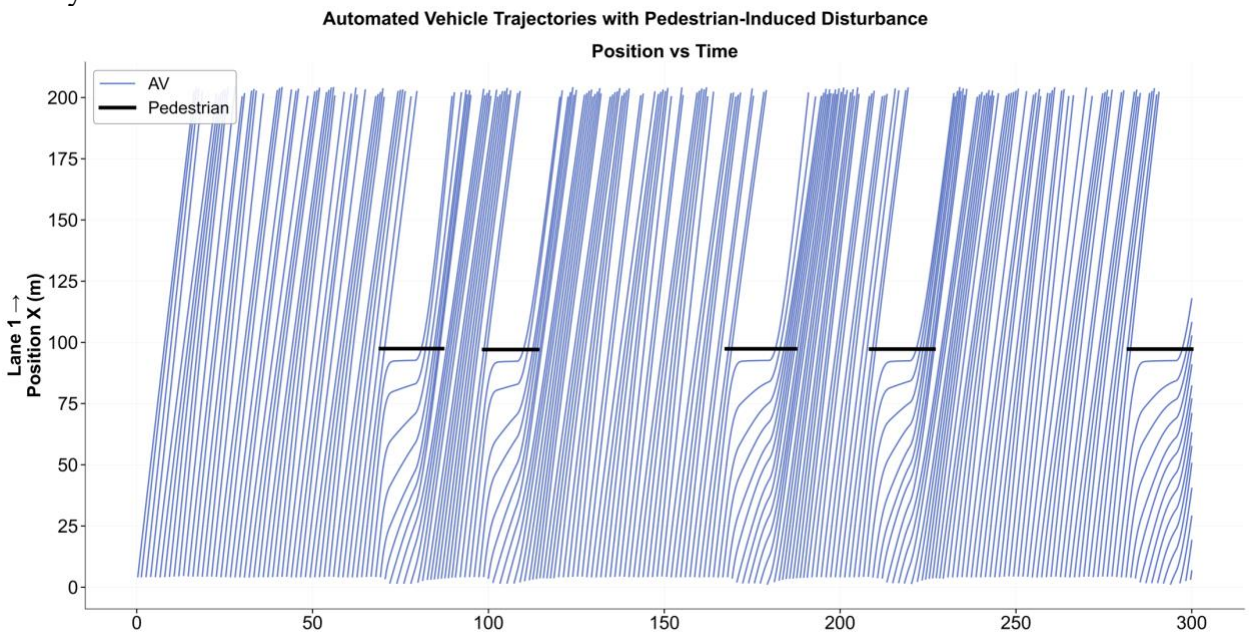

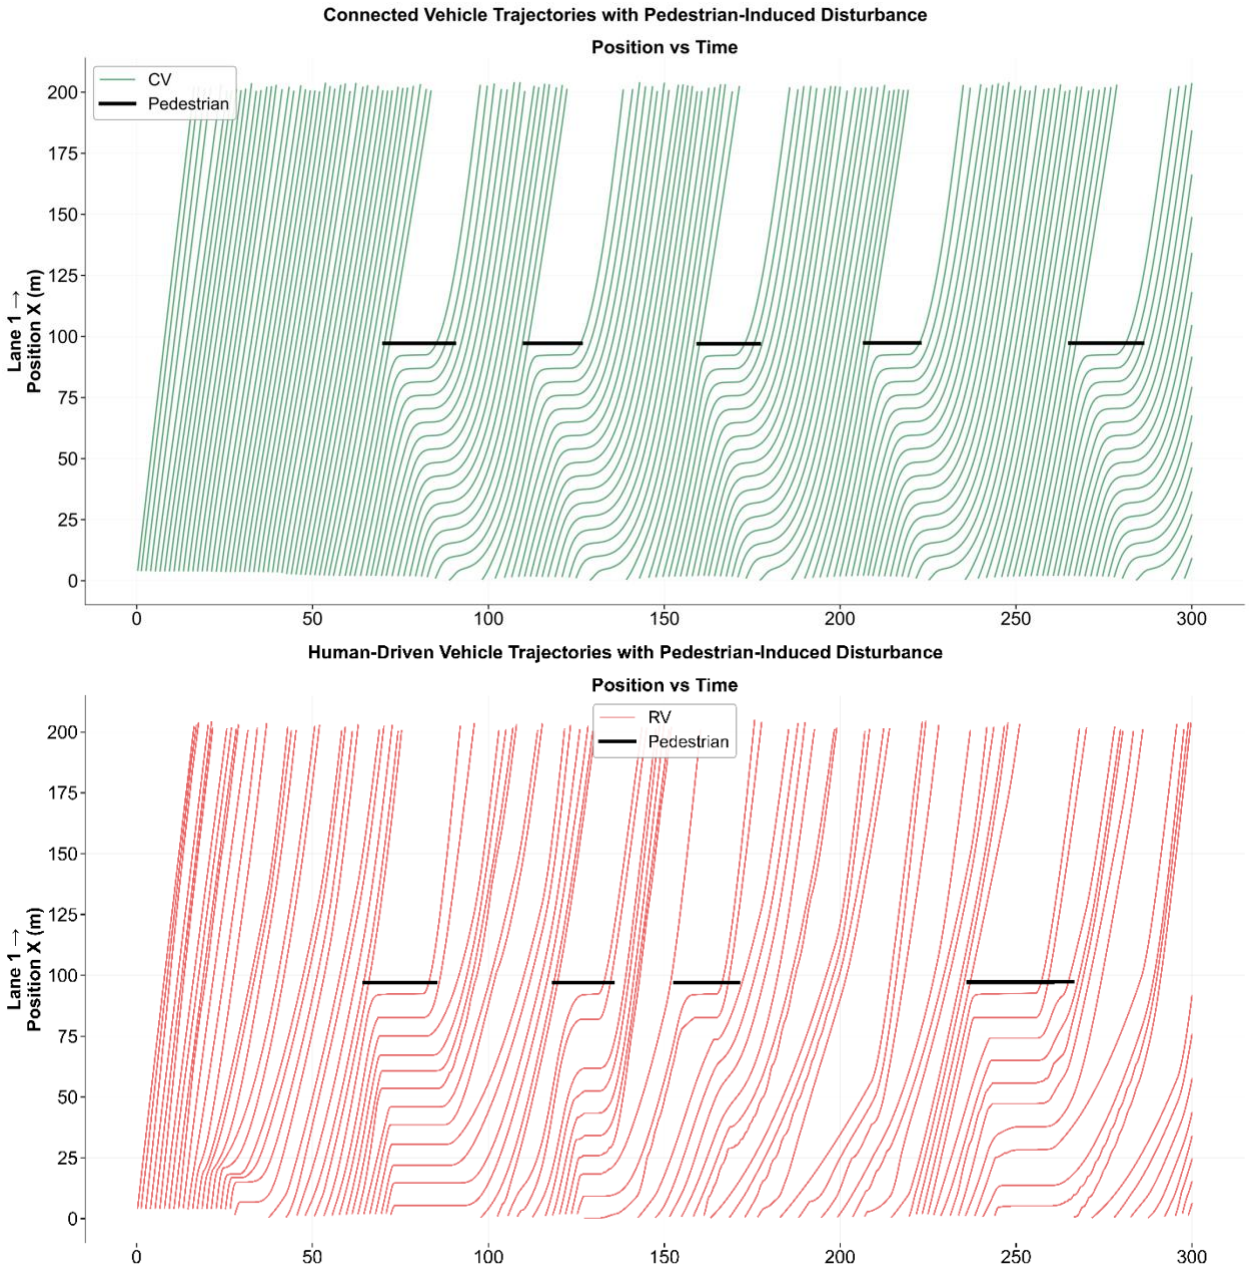

Figure 3.5 Trajectory Plots of Traffic Stream Recovery

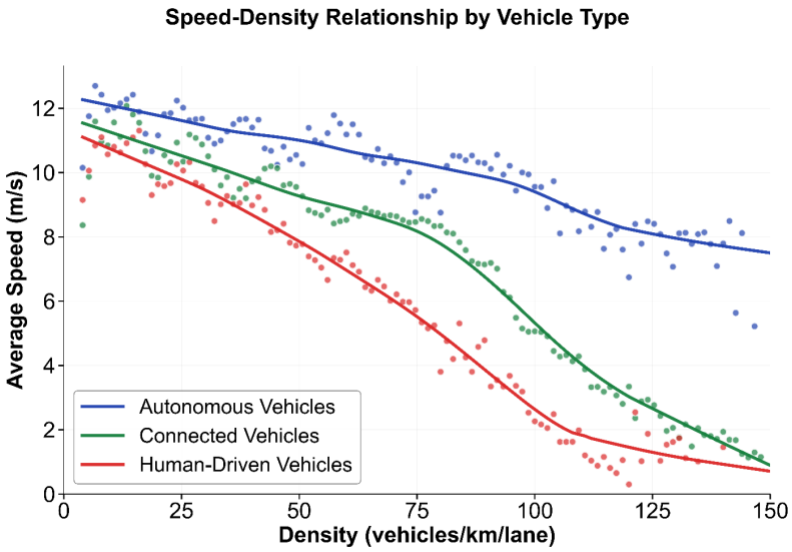

Figure 3.6 Speed Density Relationships

4 CONFLICT CHARACTERIZATION AND SEVERITY ANALYSIS

4.1 Introduction

The previous chapter described how agents navigate the street environment: pedestrians weaving through other pedestrians, cyclists negotiating speed and stability, and vehicles accelerating, braking, changing lanes, and yielding to others roadway users. These behavioral models generate agent trajectories based on local observations of neighbors. But even with realistic movement models, conflicts are not avoided altogether. Limited perception ranges, heterogeneous reaction times, and imperfect predictions of others' intentions mean that agents will inevitably encounter situations where paths intersect in unsafe ways. Understanding how these conflicts emerge, and what makes some far more dangerous than others, reveals how street design directly shapes safety outcomes.

Crucially, a conflict between two vehicles does not pose the same level of risk as one involving a pedestrian, even if the spatial arrangement and time to collision are identical. This motivates a dual-metric framework: Time to Collision (TTC) captures the temporal urgency of an interaction, while Impact Severity (IS) captures the potential consequence if the conflict were to result in collision. Weighting conflicts by both metrics distinguishes designs that merely reduce the number of interactions from those that eliminate severe ones.

4.2 Conflict Detection Methodology

4.2.1 Time-To-Collision (TTC) Conflict Identification

Physical proximity alone does not define a conflict. Relative speed, direction, timing, and agent type all shape whether an interaction is dangerous. A fast-moving vehicle approaching a cyclist head-on represents a more critical interaction than two pedestrians passing close together; likewise, interactions involving vulnerable road users generally carry greater risk than those between similarly protected agents. To systematically identify when agents are on paths that may intersect dangerously, this study uses a vectorized Time-to-Collision (TTC) calculation at each simulation step. For every pair of agents on the same link, TTC estimates the time remaining until collision if both agents maintain their current velocities. The TTC between agents i and j is defined as:

$$\mathrm{GTTC_{ij}} = \frac{d_{ij,\mathrm{adjusted}}}{v_{\mathrm{approach}}} \qquad \text{(Eq 4.1)}$$

Here, $d_{ij,\mathrm{adjusted}}$ is the separation distance adjusted for agent dimensions, and v_{approach} is the relative velocity projected onto the line connecting the two agents. Diverging or negative TTC values are discarded. A conflict is recorded when TTC is positive, below a 3-second threshold, and the pair is on a converging trajectory.

4.2.2 Conflict Severity Determination

TTC identifies when an agent is facing a situation requiring an evasive action, but it does not quantify the potential harm of a resulting collision. TTC is fundamentally a binary filter

where either the TTC is low enough to constitute a conflict or it is not. It cannot distinguish a minor low-speed interaction from a high-energy collision involving a vulnerable road user.

To capture the consequence dimension, an Impact Severity Metric is introduced, grounded in collision physics and modeled against well-established injury risk benchmarks in the literature. This metric is designed for comparative evaluation and system planning, not precise injury prediction. Its purpose is to reflect the non-linear increase in harm with speed and mass asymmetry and to provide a normalized, dimensionless measure for comparing conflict severity across agent types.

For a hypothetical two-body perfectly inelastic collision, total impact energy is:

$$E = \frac{1}{2}\mu v_{\text{rel}}^2 \text{ , } \mu = \frac{m_1 m_2}{m_1 + m_2} \quad \text{(Eq 4.2)}$$

where $v_{\text{rel}} = || \text{ v}_1 - \text{v}_2 \text{ } ||$ is is the relative speed. Energy is allocated to each agent proportional to the other agent's mass:

$$E_2 = E \times \frac{m_1}{m_1 + m_2}, E_1 = E \times \frac{m_2}{m_1 + m_2} \quad \text{(Eq 4.3)}$$

Thus, when a heavy vehicle collides with a pedestrian or cyclist, the lighter agent absorbs a disproportionate share of the impact energy, aligning with empirical vulnerability patterns.

Each agent's severity is defined as:

$$S_i = \frac{E_i}{E_{\text{threshold},i}} \quad \text{(Eq 4.4)}$$

Where $E_{\text{threshold},i}$ is a policy-based reference energy representing the onset of moderate to high injury risk for that agent type. The thresholds are calibrated using standardized masses and collision speeds drawn from widely used safety research, not individualized crash modeling. They serve as consistent weighting factors that allow cross-scenario comparisons.

Table 4.1: Agent Severity Parameters

Agent Type	Mass (Kg)	Vehicle Impact Speed	$\boldsymbol{E}_{\text{Threshold},\boldsymbol{i}}$
Pedestrian	75	35 km/hr	3000
Cyclist	90	45 km/hr	6000
Vehicle	1,500	65 km/hr	60000

These thresholds are conservative, standardized policy weights chosen to ensure the metric reliably flags interactions with meaningful safety implications.

Figure 4.1 plots severity as a function of relative speed for each collision pairing. All curves exhibit the expected quadratic shape. The horizontal line at $S = 1$ marks the high-risk severity threshold, showing where each collision type crosses into dangerous territory. The figure

highlights the dramatic sensitivity of pedestrians and cyclists to even modest impact speeds relative to vehicle–vehicle collisions.

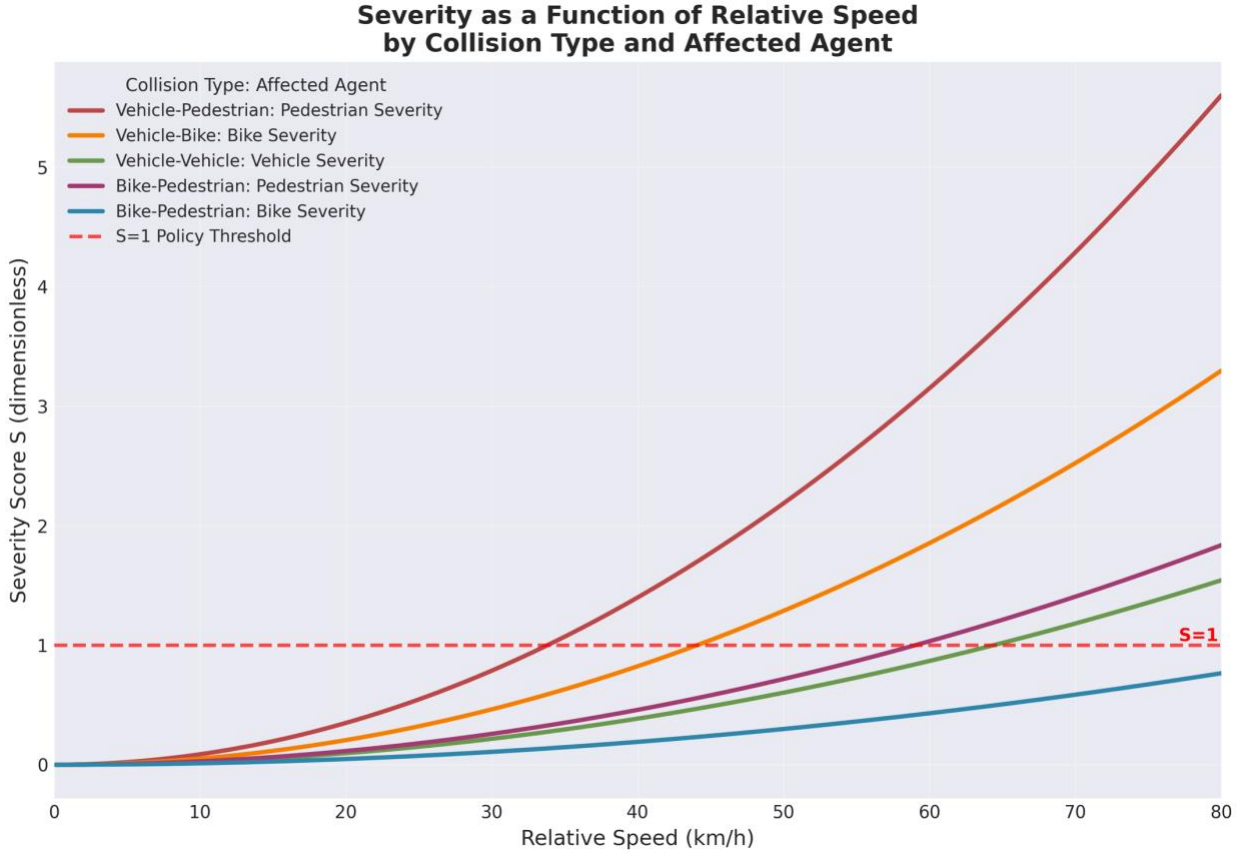

Figure 4.1 Severity Function

4.2.2.1 Brief Limitations

Although the severity metric captures the major differences across collision types, several simplifying assumptions limit its precision.

First, this model assumes perfectly inelastic collisions. The model treats every collision as if the two agents stick together and transfer the maximum possible amount of energy. Real-world impacts often involve partial elasticity as vehicles rebound, cyclists glance off the hood, pedestrians roll or slide. This assumption tends to overestimate severity for vehicle–vehicle crashes, where energy is often absorbed by crumple zones or partially returned through rebound. It may misrepresent cases where an agent is deflected rather than fully struck, such as a pedestrian being brushed and spun rather than directly hit. In other words, the inelastic assumption gives a conservative upper-bound estimate of energy transfer but does not account for variations in collision mechanics.

Second, the thresholds represent policy-level risk, not medical injury modeling. The calibrated energy thresholds correspond to transitions into “moderate-to-high risk” zones, not

clinical injury probabilities or detailed biomechanics. They are intended for comparative safety evaluation across scenarios.

Finally, angle of impact and secondary impacts like falling to the ground, sliding, or striking another object are excluded.

Despite these simplifications, the metric provides a robust and consistent basis for comparing relative conflict severity across different street designs, agent types, and scenarios.

5 THE URBAN STREET DESIGN PROBLEM

5.1 Introduction

Urban street design can be framed as a multi-objective optimization problem in which planners must balance competing goals within a highly constrained right-of-way. Two objectives are central to this trade-off. The first is throughput, reflecting how efficiently people and vehicles move through the corridor. The second is safety, operationalized as minimizing conflicts and associated risk among street users. Improvements to one objective often come at the expense of the other, making optimal design inherently dependent on how performance is evaluated. Critically, performance is not a fixed property of a street layout. Instead, it is a function of demand, which is inherently stochastic and varies by time of day, day of week, season, and surrounding land uses. As a result, evaluating street designs under idealized or deterministic demand assumptions risks misrepresenting real-world outcomes. To meaningfully compare alternative designs, it is therefore necessary to first characterize what actual demand looks like under realistic operating conditions.

5.1.1 Case Study Corridor: Belmont Avenue

To ground this analysis, a 200 m segment of Belmont Avenue in Chicago is used as a case study. This corridor is relatively short but highly active, making it well suited for detailed observation and simulation. The segment lies adjacent to the CTA Red, Brown, and Purple Line station and is surrounded by a dense mix of land uses, including multiple coffee shops, retail stores such as Target and Walgreens, restaurants, a hotel, and significant curbside activity. Together, these features generate substantial pedestrian, vehicle, and delivery traffic throughout the day. Figure 5.1 outlines this study area.

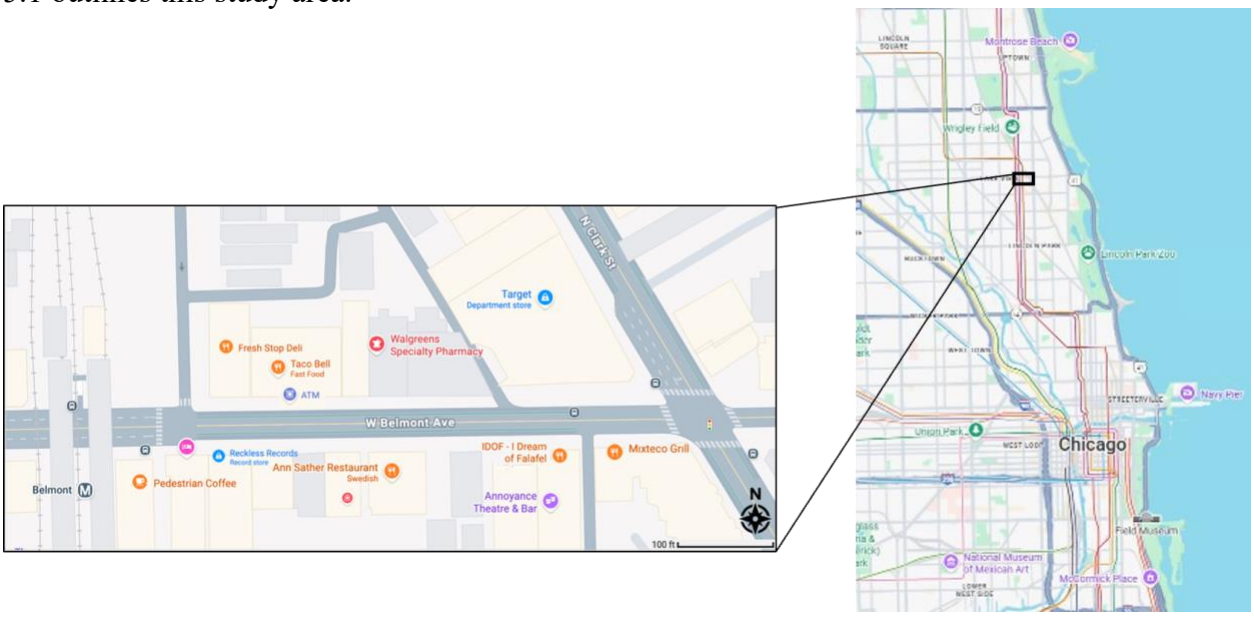

Figure 5.1 Belmont Corridor

Field observations were conducted on a warm Saturday in November (approximately 65 °F), when outdoor activity levels were relatively high. Although demand increases further during

the summer months, activity during the observation period was already substantial. In a single hour, 12 pedestrians were observed jaywalking mid-block to access destinations such as Walgreens and Taco Bell, despite the presence of marked crossings at both ends of the 200 m corridor. Pedestrian volumes during this period were nearly comparable to vehicle volumes, and nontraditional users, including a small autonomous delivery robot, were also present on the sidewalk. These observations highlight the diversity of agents and the frequency of informal, unscheduled interactions that shape safety and flow on urban streets.

Table 5.1 Field Counts

Mode / Activity	1st 15 min	2nd 15 min	3rd 15 min	4th 15 min
Cars (Eastbound)	190	147	155	210
Cars (Westbound)	185	192	198	205
Buses (Eastbound)	2	3	1	3
Buses (Westbound)	2	1	3	2
Pedestrians (Eastbound)	144	152	130	103
Pedestrians (Westbound)	148	192	107	103
Bicycles	6	7	9	8
E-bikes	2	0	1	1
E-scooters	1	2	1	1
Trucks	1	3	2	1
Trolley	1	0	1	0
Delivery robot	0	1	0	0
Jaywalking pedestrians	3	2	4	3
Rideshare pick-up / drop-off	6	5	9	6

5.1.2 Demand Characterization and Simulation Framework

Using these field counts, a baseline demand schedule was constructed that distinguishes between flowing agents and stopping agents. Flowing agents include pedestrians, cyclists, and vehicles whose primary behavior is continuous movement through the corridor, characterized by desired speeds and realized flow rates. Stopping agents include pedestrians crossing mid-block, curbside deliveries, and other activities that temporarily disrupt traffic flow, characterized by the timing and location of stops or crossings.

Because real-world demand is variable, this schedule represents only one possible realization of activity on the corridor. To account for stochasticity, the simulation is run multiple times using a Monte Carlo approach. For each run, flowing agents' desired speeds and arrival rates are drawn from probability distributions, while stopping agents' crossing times and locations are randomly perturbed within ±2 minutes and ±10 meters of their observed values. Each simulation run thus represents a distinct but plausible realization of demand.

Simulations are executed in parallel, and performance metrics are aggregated across runs to estimate expected outcomes. The framework overall is outlined in the flowchart shown in Figure

5.2.

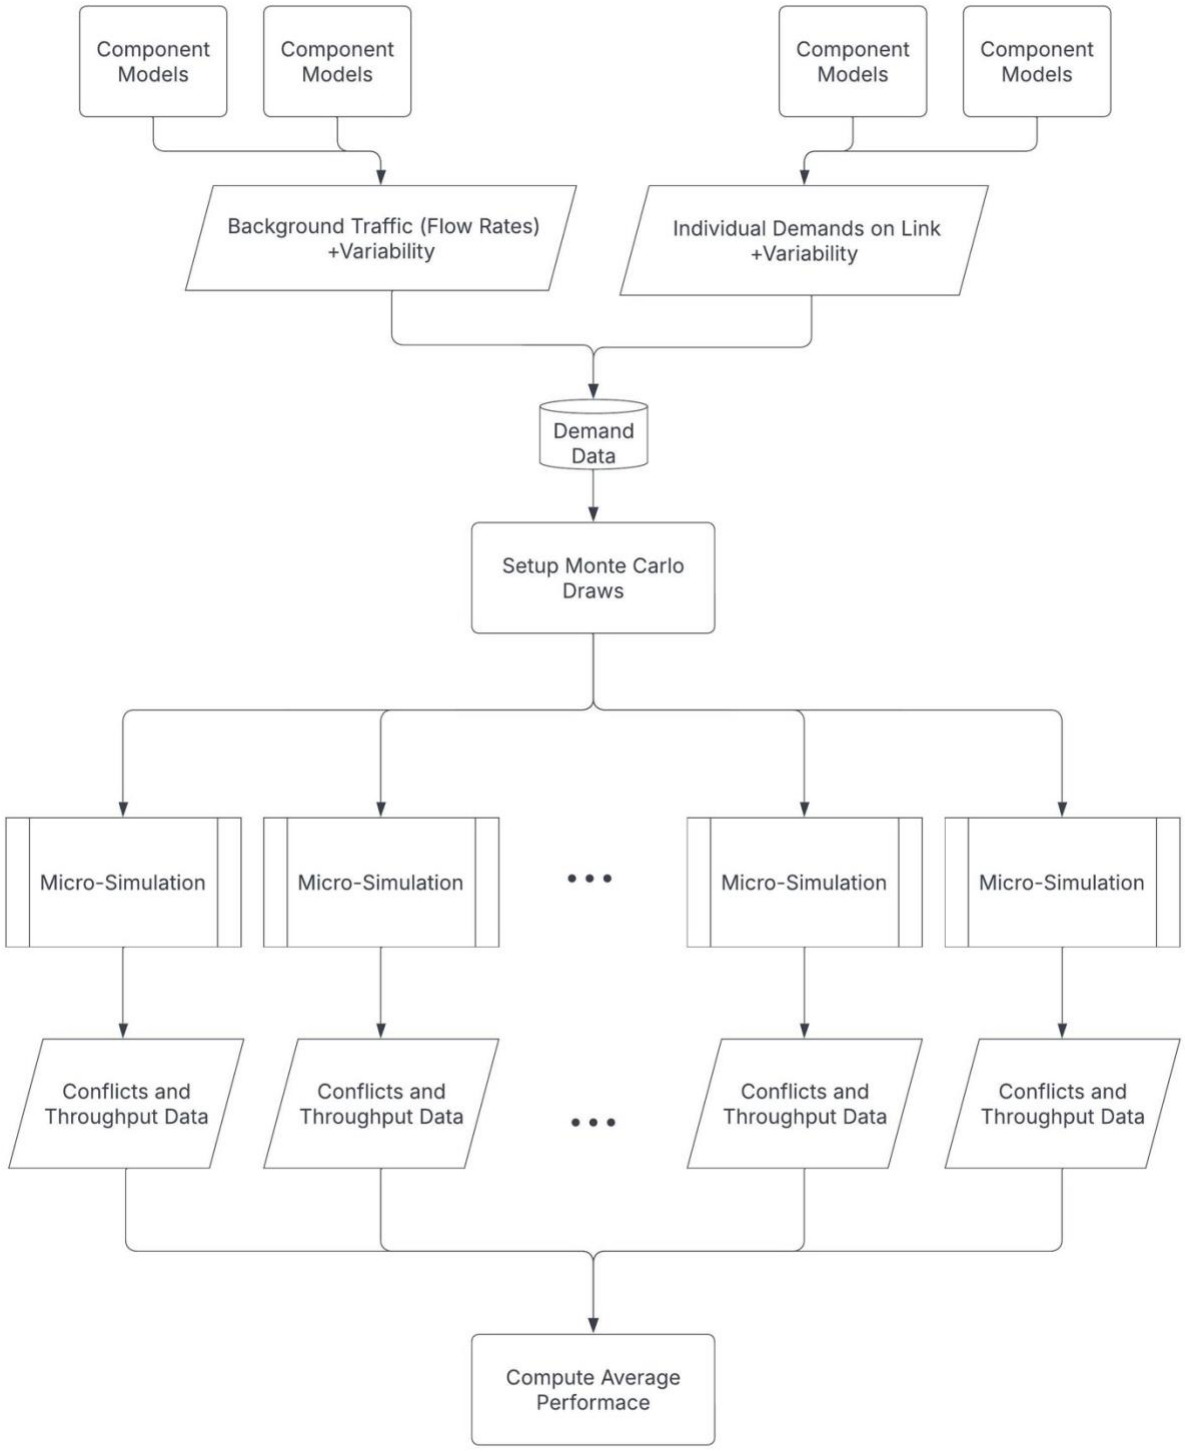

Figure 5.2 Framework Flowchart

In practice, results stabilize after approximately 40 simulation runs, indicating convergence of the estimated performance measures. This process yields robust estimates of both throughput

and conflict outcomes for a given street design under realistic operating conditions. Figure 5.3 shows this convergence plot.

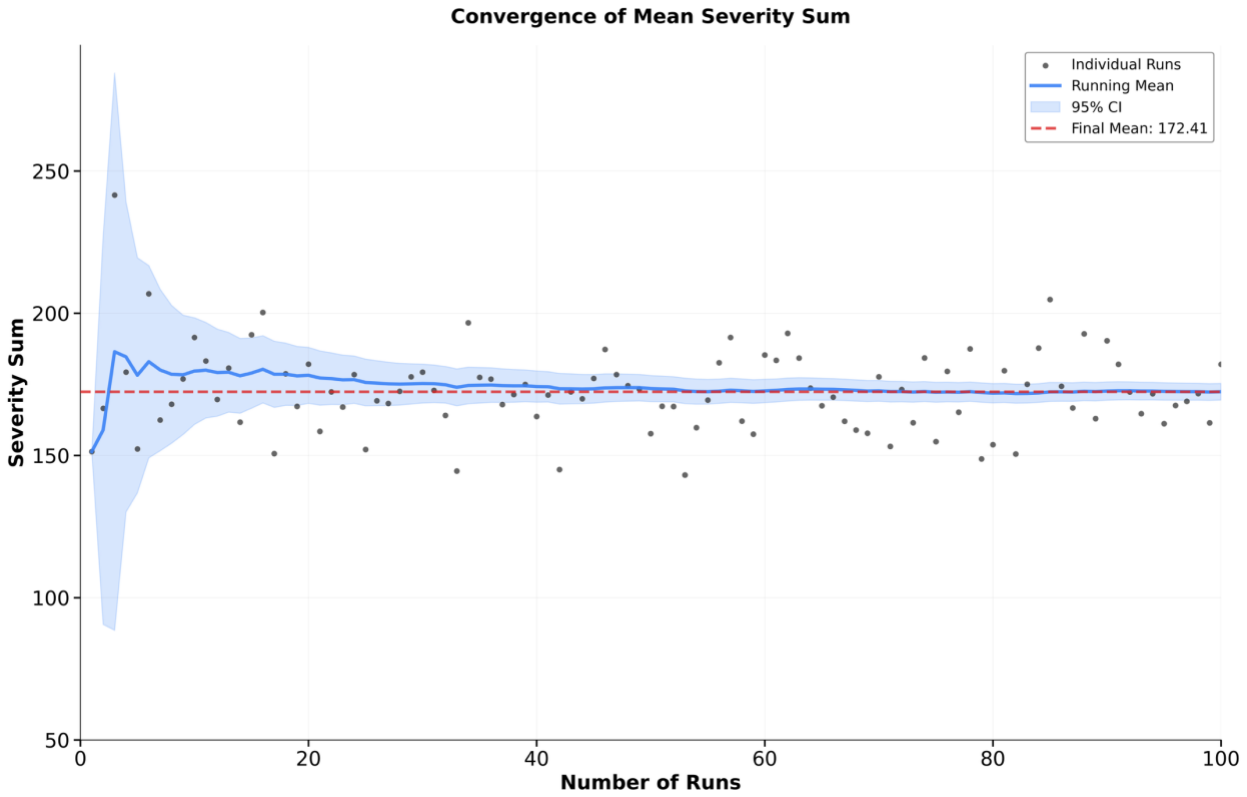

Figure 5.3 Convergence of Severity Score

5.1.3 Feasible Street Designs

To systematically explore alternative street layouts, this study defines a bounded set of feasible street designs based on geometric and operational constraints that reflect common urban design standards. The goal is not to prescribe a single ideal configuration, but rather to identify the full range of layouts that can physically fit within the available right-of-way while supporting bidirectional travel and multimodal use.

All candidate designs are required to provide at least one vehicle lane in each direction, ensuring continuous eastbound and westbound traffic flow. Let L_{east} and L_{west} denote the number of eastbound and westbound vehicle lanes, respectively; both must be positive integers. Each vehicle lane is assigned a uniform width, denoted by w_{veh}, which is constrained to fall between 10 and 15 feet to remain consistent with standard lane design guidance.

Pedestrian infrastructure is explicitly represented through sidewalks on both sides of the street. The width of each sidewalk, denoted by $w_{sidewalk}$, must be at least 3 feet. While this minimum does not imply ideal pedestrian conditions, it establishes a lower bound that ensures basic pedestrian accommodation in all designs.

Bicycle infrastructure is treated as optional. A binary indicator variable, z_{bike}, is used to represent the presence or absence of dedicated bike lanes. When $z_{\text{bike}} = 1$, a bike lane is provided on each side of the street, with each lane assigned a width w_{bike} between 3 and 8 feet. When $z_{\text{bike}} = 0$, no dedicated bike lanes are included and the bike lane width is set to zero.

All design elements must collectively fit within the available cross-sectional roadway width, denoted by W_{avail}. This total width constraint ensures that sidewalks, vehicle lanes, and bike lanes together do not exceed the physical right-of-way. Formally, the feasibility of a street design is enforced through the following constraint:

$$2\, w_{\text{sidewalk}} + (L_{\text{east}} + L_{\text{west}})\, w_{\text{veh}} + 2\, z_{\text{bike}}\, w_{\text{bike}} \;\leq\; W_{\text{avail}} \qquad \text{(Eq 5.1)}$$

Each unique combination of lane counts, lane widths, sidewalk widths, and bike lane configurations that satisfies this constraint constitutes a distinct street design. Applying these rules to the 60-foot right-of-way considered in the Belmont Avenue case study produces nearly 200 feasible street configurations. This finite but diverse design space captures a wide spectrum of trade-offs between pedestrian space, vehicle capacity, and bicycle accommodation, and serves as the basis for the simulation and optimization analyses that follow.

5.1.4 Enumerating and Evaluating Street Designs

Once demand is fixed, the next step is to vary the street design itself and repeat the evaluation process. For the Belmont Avenue case study, the available right-of-way is approximately 60 feet, which significantly constrains feasible layouts. Design constraints include practical bounds on lane widths: vehicle lanes must generally fall between 10 and 15 feet, while bicycle lanes range from approximately 3 to 8 feet. Within these constraints, there are on the order of 200 distinct feasible layouts.

Because the solution space is relatively small in this case, all feasible designs can be explicitly enumerated rather than searched using heuristics. Each layout is evaluated using the same stochastic demand framework, producing paired estimates of throughput and safety performance. Figure 5.4 shows the design space and how various layout performed in relation to design variables.

Figure 5.4 Design Space

5.1.5 Pareto Frontier and Design Trade-offs

With performance measured along two objectives, results can be visualized using a Pareto frontier. Each point on the frontier corresponds to a non-dominated design, meaning that no other layout performs better on both throughput and safety simultaneously. Some designs prioritize safety at the cost of reduced flow, while others maximize throughput while accepting higher conflict levels.

A scalar weighting parameter, α, is used to trace this frontier. When $\alpha = 0$, the optimization prioritizes safety alone; when $\alpha = 1$, it prioritizes throughput alone. Intermediate values of α yield compromise solutions that balance the two objectives. For any chosen value of α, a single optimal design can be identified, corresponding to the point labeled on the frontier.

Although the Belmont Avenue case study involves a limited number of design alternatives, the framework readily generalizes to more complex environments. As the number of streets, design variables, or objectives increases, the solution space grows rapidly, motivating the need for

scalable simulation and optimization methods. Nonetheless, this case study demonstrates how realistic demand modeling, stochastic simulation, and Pareto-based evaluation can be combined to rigorously assess trade-offs in urban street design.

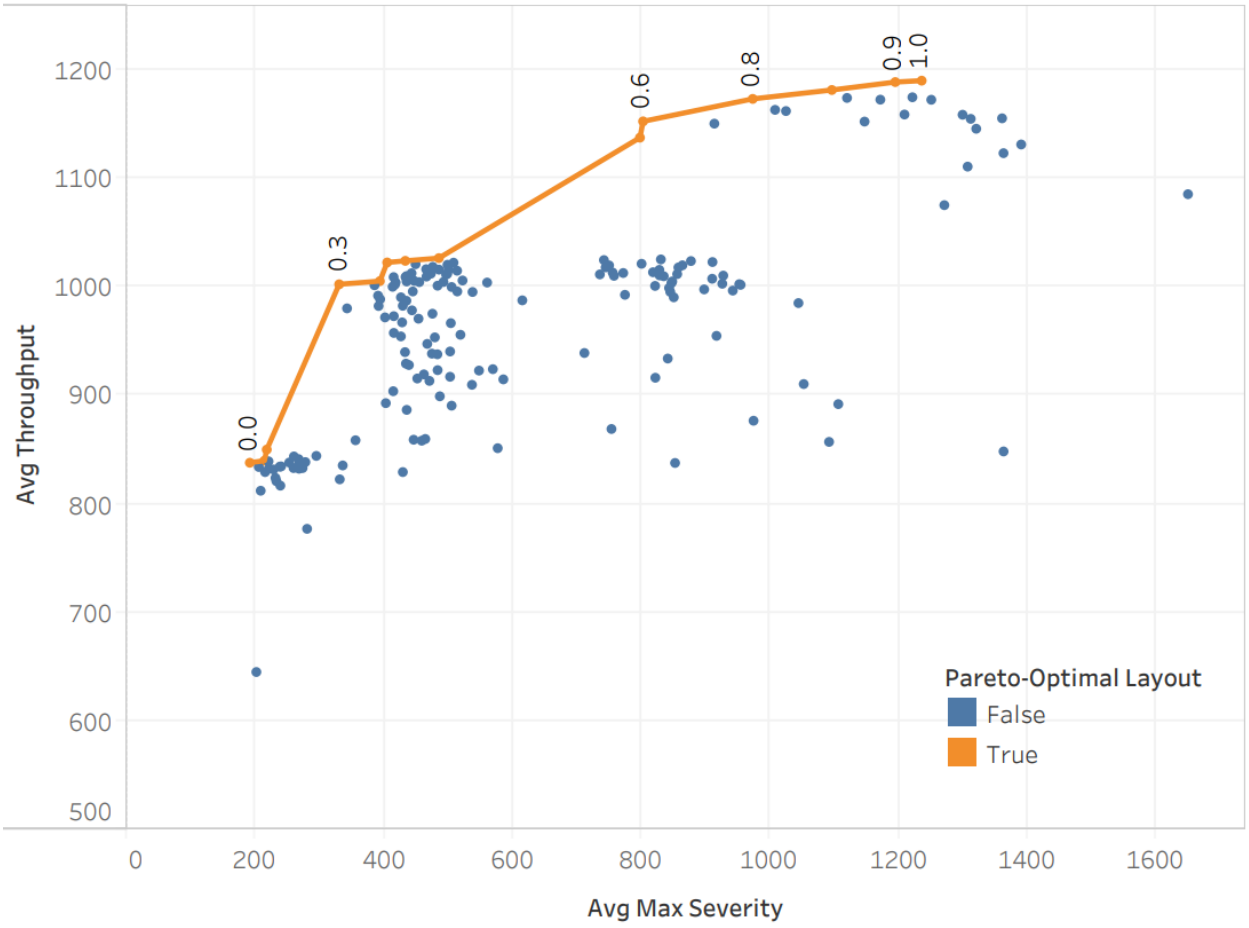

Figure 5.5 Pareto Frontier of Designs

5.1.6 Conclusions

The conflict characterization and severity analysis framework presented in this chapter provides a powerful tool for evaluating the safety implications of urban street designs under realistic operating conditions. By combining Time-to-Collision (TTC) metrics with an impact severity measure grounded in collision physics, the approach moves beyond simple conflict counts to capture both the likelihood and potential consequences of interactions among diverse street users. This dual-metric system enables planners to distinguish between designs that merely reduce the number of conflicts and those that meaningfully mitigate severe outcomes for vulnerable road users such as pedestrians and cyclists.

These tools can be applied to compare alternative layouts, assess trade-offs between throughput and safety, and identify design elements, such as protected bike lanes or reduced shared space, that significantly lower high-risk interactions. However, while the framework provides a rigorous quantitative basis for evaluation, its application must be guided by policy

priorities and community values. Decisions about how to balance efficiency, equity, and safety cannot be resolved by technical metrics alone; they require input from stakeholders and alignment with broader goals such as Vision Zero, sustainability, and accessibility. In this sense, the methodology serves as an evidence-based decision-support system rather than a prescriptive solution, enabling transparent discussions about trade-offs and helping communities select designs that reflect their priorities.

Findings

The research demonstrates that integrated modeling and optimization tools can effectively support urban street design decisions that balance efficiency, safety, and equity. The two-step optimization framework based on the Network Fundamental Diagram (NFD) successfully identifies flow-maximizing network attributes and translates them into practical link-level configurations using a Genetic Algorithm (GA). While the GA cannot fully replicate the theoretical optimal NFD due to physical and budget constraints, it consistently improves accessibility and reduces conflict potential compared to the existing network.

Across all formulations, utilitarian, sufficiency, accessibility gap, and maximin, the optimized networks outperform the baseline in terms of multimodal connectivity. Bike accessibility within a 20-minute threshold increases substantially, and equity metrics such as access for low-income zones show modest improvements. Ethical objectives influence spatial allocation patterns:

- Utilitarian designs prioritize overall accessibility and flow efficiency.
- Sufficiency-based designs expand coverage to ensure minimum access for all zones.
- Accessibility gap and maximin formulations target fairness by reducing disparities and improving outcomes for disadvantaged groups.

Simulation-based evaluations confirm that these designs reduce severe conflicts without compromising throughput excessively. The multi-agent simulation platform captures realistic interactions among vehicles, cyclists, and pedestrians, enabling conflict characterization through Time-to-Collision (TTC) and Impact Severity metrics. Results indicate that layouts with dedicated micromobility infrastructure and reduced shared space achieve lower severity scores, particularly for vulnerable road users.

Finally, the case study on Belmont Avenue illustrates the practical feasibility of these tools. Nearly 200 candidate designs were evaluated under stochastic demand conditions, and Pareto analysis revealed clear trade-offs between safety and efficiency. Designs that allocate space for protected bike lanes and wider sidewalks consistently rank higher on safety, while layouts with more lanes maximize throughput. These findings underscore the importance of transparent trade-off analysis and community-driven prioritization in selecting final designs.

Recommendations

Urban planners and policymakers should incorporate quantitative modeling and simulation tools as a central part of street design and evaluation. Traditional qualitative guidelines provide important principles but often fail to capture the complex trade-offs between safety, efficiency, and equity in multimodal environments. The integrated frameworks presented in this study combine Network Fundamental Diagram (NFD)-based optimization, multi-agent simulation, and conflict severity analysis to deliver a rigorous, data-driven foundation for decision-making. These tools enable transparent evaluation of design alternatives, quantify accessibility and flow impacts, and assess safety outcomes under realistic operating conditions. Using such approaches will allow communities to move beyond ad hoc design practices and toward evidence-based solutions that balance technological innovation with human-centered values.

Outputs

This research delivers two primary outputs that advance the state of practice in multimodal street design and evaluation:

1. New Methodologies for Network Design Optimization

   A two-step optimization framework was developed to systematically design multimodal networks. The first step applies a Network Fundamental Diagram (NFD)-based formulation to identify flow-maximizing values of network attributes such as mode-exclusive areas, shared space, and interaction plane length. The second step translates these aggregate optimal characteristics into link-level configurations using a Genetic Algorithm (GA). This approach enables planners to balance efficiency, accessibility, and fairness by incorporating ethical principles such as sufficiency, accessibility gap, and maximin into the optimization process. These formulations move beyond traditional qualitative guidelines by providing a quantitative basis for evaluating trade-offs between throughput, equity, and safety under realistic constraints.

2. Multi-Agent Simulation Platform for Urban Streets

   A comprehensive simulation platform was created to model the dynamic interactions among diverse street users, including connected and autonomous vehicles, human-driven vehicles, bicycles, pedestrians, and emerging micromobility modes. The platform integrates microscopic behavioral models, such as car-following, lane-changing, pedestrian crossing decisions, and shared-lane yielding, with stochastic demand generation and hierarchical decision-making. This enables realistic representation of both flowing and stopping agents and captures emergent macroscopic phenomena such as shockwaves and congestion patterns. The simulator supports conflict detection using Time-to-Collision (TTC) and Impact Severity metrics, allowing quantitative assessment of safety outcomes across alternative street designs. By combining behavioral realism with performance evaluation, this tool provides planners with actionable insights into how design choices influence efficiency and safety in complex urban environments.

Together, these outputs establish a rigorous, data-driven foundation for designing and evaluating multimodal streets. They enable transparent comparison of design alternatives, quantify accessibility and flow impacts, and assess safety under stochastic demand conditions, supporting evidence-based decision-making for future urban mobility.

Outcomes

The application of the developed methodologies and simulation platform has produced several important outcomes:

1. Increased Understanding and Awareness
   a. The research enhances understanding of how multimodal interactions and CAV deployments affect safety, efficiency, and equity on urban streets. It introduces quantitative tools that move beyond qualitative guidelines, enabling planners to evaluate trade-offs transparently.

2. Expansion of the Body of Knowledge
   a. The study contributes new formulations for multimodal network design optimization and introduces a multi-agent simulation platform that captures realistic interactions among diverse street users. These outputs advance the state of practice in urban mobility research.

3. Improved Processes and Techniques
   a. The two-step optimization framework and simulation platform provide planners with systematic, data-driven processes for designing multimodal networks. These processes improve decision-making by integrating flow efficiency, accessibility, and fairness objectives.

4. Adoption Potential for New Practices
   a. The tools developed in this research can be incorporated into planning workflows, supporting evidence-based design and policy decisions. They enable agencies to adopt quantitative evaluation methods for Complete Streets and CAV integration.

Impacts

The outcomes of this research have the potential to generate significant impacts on transportation systems and society:

1. Enhanced Safety for Vulnerable Road Users
    a. Designs informed by the simulation platform and optimization framework reduce severe conflicts, particularly for pedestrians and cyclists, contributing to lower crash risk and improved street safety.

2. Improved Efficiency and Accessibility
    a. Optimized network configurations increase multimodal accessibility while maintaining reasonable throughput, supporting more equitable and efficient urban mobility.

3. Support for Policy and Investment Decisions
    a. By providing transparent trade-off analysis, these tools can guide infrastructure investments and policy frameworks toward designs that balance technological innovation with human-centered values.

4. Long-Term Community and Environmental Benefits
    a. Safer, more efficient multimodal streets encourage active transportation, reduce reliance on private vehicles, and contribute to sustainability goals through lower emissions and improved livability.

Challenges and Lessons Learned

Developing large-scale optimization and simulation frameworks for multimodal street design presented several technical and methodological challenges. One of the primary difficulties was ensuring computational efficiency when handling highly complex models that integrate multiple transportation modes, stochastic demand patterns, and ethical constraints. The optimization formulations, particularly those involving nonlinear objectives and multiple constraints, required advanced solvers and careful tuning to achieve convergence within reasonable time frames. Implementing the two-step optimization process and genetic algorithm for link-level design demanded custom operators and repair functions to maintain feasibility and spatial coherence, which added to development complexity.

On the simulation side, building a multi-agent platform capable of representing diverse behaviors, such as car-following, lane-changing, pedestrian crossing, and shared-lane yielding, posed significant challenges in terms of scalability and realism. Modeling interactions among thousands of agents while preserving detailed behavioral rules required hierarchical decision structures and efficient data handling. Additionally, integrating conflict detection and severity analysis into the simulation workflow introduced further computational overhead, necessitating optimizations in both algorithm design and software architecture.

These challenges underscore the importance of balancing model fidelity with computational tractability. Lessons learned include the need for modular software design, parallel processing for simulation runs, and adaptive optimization strategies to handle non-convex problems. Future work should focus on improving solver performance, leveraging high-performance computing resources, and refining behavioral models to reduce complexity without sacrificing accuracy.

References

[1] T. Dai, J. Li, and M. Nie, "Accessibility-Based Ethics-Aware Transit Design," Nov. 04, 2022, *Rochester, NY*: 4268432. doi: 10.2139/ssrn.4268432.

[2] M. B. Lowry, P. Furth, and T. Hadden-Loh, "Prioritizing new bicycle facilities to improve low-stress network connectivity," *Transp. Res. Part Policy Pract.*, vol. 86, pp. 124–140, Apr. 2016, doi: 10.1016/j.tra.2016.02.003.

[3] City of Evanston, "CHICAGO AVENUE Multimodal Corridor Improvement," Feb. 2023.

[4] "The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics | IEEE Journals & Magazine | IEEE Xplore." Accessed: Dec. 20, 2025. [Online]. Available: https://ieeexplore.ieee.org/document/4019451

[5] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.*, vol. 368, no. 1928, pp. 4585–4605, 2010.

[6] S. H. Hamdar, M. Treiber, H. S. Mahmassani, and A. Kesting, "Modeling Driver Behavior as Sequential Risk-Taking Task," *Transp. Res. Rec.*, vol. 2088, no. 1, pp. 208–217, Jan. 2008, doi: 10.3141/2088-22.

[7] A. Talebpour, H. S. Mahmassani, and S. H. Hamdar, "Multiregime Sequential Risk-Taking Model of Car-Following Behavior: Specification, Calibration, and Sensitivity Analysis," *Transp. Res. Rec.*, vol. 2260, no. 1, pp. 60–66, Jan. 2011, doi: 10.3141/2260-07.

[8] M. Haque, S. H. Hamdar, and A. Talebpour, "Bridging the gap between micro-economics and micro-mobility: A two-dimensional risk-based microscopic model of pedestrians' and bicyclists' operational behaviors," *Transp. Res. Part B Methodol.*, p. 103021, July 2024, doi: 10.1016/j.trb.2024.103021.

[9] A. Kesting, M. Treiber, and D. Helbing, "General Lane-Changing Model MOBIL for Car-Following Models," *Transp. Res. Rec.*, vol. 1999, no. 1, pp. 86–94, Jan. 2007, doi: 10.3141/1999-10.

[10] A. Talebpour, H. S. Mahmassani, S. H. Hamdar, and University of Illinois at Urbana-Champaign, "Third Generation Simulation Data (TGSIM): A Closer Look at The Impacts of Automated Driving Systems on Human Behavior," FHWA-JPO-24-133, May 2024. Accessed: June 08, 2024. [Online]. Available: https://rosap.ntl.bts.gov/view/dot/74647

[11] A. Khakpour, S. J. Hegde, N. Li, and H. S. Mahmassani, "A Drift Diffusion Model of Ped-Crossing Interactions with AVs on Two-lane Complete Streets," presented at the TRB 104th Annual Meeting, Jan. 2025.