THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

École doctorale : EEATS - Electronique, Electrotechnique, Automatique, Traitement du Signal
Spécialité : Automatique - Productique
Unité de recherche : Centre de recherche Inria de l'Université Grenoble Alpes

Conception optimale du réseau de mobilité urbaine pour un partage durable de l'espace entre les véhicules et les modes de transport doux

Optimal urban mobility network design for sustainable space sharing between vehicles and soft transport modes

Présentée par :

Manuel CAMPERO JURADO

Direction de thèse :

Carlos CANUDAS DE WIT DIRECTEUR DE RECHERCHE, CNRS	Directeur de thèse
Giovanni DE NUNZIO IFP Energies nouvelles	Co-encadrant de thèse

Rapporteurs :

Rich JEPPE
FULL PROFESSOR, Technical University of Denmark
Luca CONSOLINI
ASSOCIATE PROFESSOR, University of Parma

Thèse soutenue publiquement le **13 mars 2026**, devant le jury composé de :

Sophie ACHARD, DIRECTRICE DE RECHERCHE, CNRS délégation Alpes	Présidente
Rich JEPPE, FULL PROFESSOR, Technical University of Denmark	Rapporteur
Luca CONSOLINI, ASSOCIATE PROFESSOR, University of Parma	Rapporteur
Alain KIBANGOU, MAITRE DE CONFERENCES HDR, Université Grenoble Alpes	Examinateur
Thomas BUHLER, PROFESSEUR DES UNIVERSITES, Université Marie et Louis Pasteur	Examinateur

Invités :

Mauro SALAZAR
ASSOCIATE PROFESSOR, Eindhoven University of Technology

UNIVERSITÉ DE GRENOBLE ALPES

ÉCOLE DOCTORALE EEATS

Eléctronique, Eléctrotechnique, Automatique, Traitement du Signal

T H È S E

pour obtenir le titre de

docteur en sciences

de l'Université de Grenoble

Mention : Automatique et Productique

Présentée et soutenue par

Manuel CAMPERO JURADO

Optimal urban mobility network design for sustainable space sharing between vehicles and soft transport modes

Thèse dirigée par **Carlos CANUDAS-DE-WIT**
co-encadrée par **Giovanni DE NUNZIO**

préparée au laboratoire Grenoble Images Parole Signal Automatique (GIPSA-Lab)

soutenue le 13/03/2026

Jury :

Rapporteurs :	Jeppe RICH	-	Technical University of Denmark, Denmark
	Luca CONSOLINI	-	University of Parma, Italy
Présidente :	Sophie ACHARD	-	Univ. Grenoble Alpes, CNRS, INRIA Grenoble INP, LJK, France
Examinateurs :	Alain KIBANGOU	-	Université Grenoble Alpes, Gipsa-Lab, France
	Thomas BUHLER	-	Université Marie et Louis Pasteur, France
Invité :	Mauro SALAZAR	-	Eindhoven University of Technology, The Netherlands

Résumé — Le vélo peut jouer un rôle majeur pour rendre les villes plus saines, plus propres et plus soutenables sur le plan social, mais son usage quotidien reste limité dans de nombreux contextes, car beaucoup de personnes perçoivent la circulation au contact du trafic motorisé comme dangereuse ou inconfortable. Comme le degré de séparation entre cyclistes et véhicules motorisés est un déterminant essentiel de la sécurité, perçue comme objective, cette thèse formule l'amélioration de la sécurité à vélo comme un problème de conception de réseau sous contrainte budgétaire. Le réseau cyclable est modélisé comme un graphe orienté pondéré, dont les liens reçoivent un degré de séparation (DoS) à six niveaux dérivé d'OpenStreetMap, fournissant des scores portables et orientés sécurité. En s'appuyant sur l'idée que la sécurité est aussi une propriété du réseau, façonnée par la connectivité et les options d'itinéraires, nous identifions d'abord des métriques de théorie des graphes qui mesurent si les aménagements sûrs forment des corridors continus plutôt que des segments fragmentés. Nous étudions ensuite le comportement monotone de ces métriques sous des améliorations de DoS afin de garantir que des améliorations d'infrastructure ne puissent pas dégrader l'objectif. Ces métriques sont reliées à des données d'enquête fondées sur les perceptions pour construire un modèle prédictif, puis utilisées dans une optimisation des décisions d'aménagement sous budget. Plutôt que de supposer que la topologie suffit à hiérarchiser les améliorations, nous illustrons, par un exemple simple, qu'une lecture purement structurelle peut mal allouer l'investissement lorsque des liens centraux portent une faible demande cycliste. Cela motive une extension sensible aux flux, où les déplacements sont représentés par une affectation all-or-nothing fondée sur une distance généralisée spécifique au vélo, tenant compte de la pente, des changements angulaires et de l'exposition au trafic. On obtient ainsi une mesure d'exposition pondérée par le nombre de cyclistes concernés, avec une rétroaction bi-niveau : les améliorations modifient les coûts d'itinéraires, les coûts redistribuent les flux, et les flux déterminent l'exposition finale. Le cadre est ensuite étendu au multimodal pour rendre explicite le compromis lors d'une réallocation de l'espace viaire : réduire l'exposition des cyclistes peut augmenter les temps de trajet en voiture via des changements de capacité. Dans cette continuité, nous intégrons le report modal induit et des objectifs de justice des transports en permettant aux améliorations cyclables d'influencer le choix modal, et en comparant des objectifs utilitaristes, visant la performance globale, à des objectifs suffisantaristes, visant à réduire la part d'usagers sous des seuils minimaux de sécurité ou d'accessibilité. Enfin, pour comparer les villes, nous proposons une réduction de réseau préservant l'exposition, qui regroupe de grands graphes cyclables en super-réseaux de taille fixe. Cela permet un benchmarking plus équitable malgré les différences de taille et de résolution cartographique, et soutient des recommandations plus interprétables et plus transférables selon la structure des réseaux.

Mots clés : Infrastructures cyclables, conception de réseaux, métriques de théorie des graphes, sécurité des cyclistes, trafic multi-modal, justice des transports, mobilités urbaines, agrégation par clustering.

Abstract — Cycling can play a major role in making cities healthier, cleaner, and more socially sustainable, yet everyday cycling remains limited in many places because many people perceive riding alongside motor traffic as unsafe or uncomfortable. Since the degree of separation between cyclists and motorized vehicles is a key determinant of both perceived and objective safety, this thesis formulates cycling-safety improvement as a budget-constrained network design problem. The cycling network is modeled as a weighted directed graph whose links are assigned a six-level Degree of Separation (DoS) derived from OpenStreetMap, providing portable, safety-oriented scores. Building on the idea that safety is not only a property of individual road segments but also a network property shaped by connectivity and route options, we first identify graph theory metrics that quantify whether safe facilities form continuous corridors rather than fragmented segments. We analyze the monotonic behavior of these metrics under DoS upgrades to ensure that infrastructure "improvements" cannot paradoxically worsen the objective. These metrics are then linked to perception-based survey data to build a predictive model and to support budgeted optimization of upgrade decisions. Instead of assuming that topology alone is sufficient for prioritizing upgrades, we illustrate with a simple example that a purely structural view can misallocate investment when centrally located links carry little cycling demand. This motivates a flow-aware extension of the framework, where cyclist movements are represented through an all-or-nothing assignment based on a cyclist-specific generalized distance that accounts for slope, angular changes, and traffic exposure. The result is an exposure metric weighted by the number of cyclists affected and a bi-level feedback mechanism in which upgrades modify route costs, costs reshape route choice and flows, and the redistributed flows determine the resulting exposure. The framework is then extended to a multi-modal setting to make explicit the trade-off faced by planners when reallocating road space: reducing cyclist exposure may increase car travel times through capacity changes. Building on this, we incorporate induced modal shift and transport-justice objectives by allowing improved cycling conditions to influence mode choice, and by comparing utilitarian objectives, which aim to improve aggregate system performance, with sufficientarian objectives, which aim to reduce the share of users below minimum safety or accessibility thresholds. This comparison highlights how distributional goals and demand feedback can substantially change which projects are selected. Finally, to support cross-city comparison, we propose an exposure-preserving network reduction that clusters large cycling graphs into fixed-size super-networks. This yields fairer benchmarking despite differences in network size and mapping resolution, and it supports more interpretable and transferable recommendations tailored to different network structures.

Keywords: Cycling infrastructure, network design, graph-theory metrics, cyclist safety, multi-modal traffic, transport justice, urban mobility, clustering.

Grenoble Images Parole Signal Automatique (GIPSA-Lab)
11 Rue des Mathématiques, 38400 Saint-Martin-d'Hères

Résumé

Le vélo est largement reconnu comme un levier pour construire des villes plus saines, plus propres et plus inclusives. Il améliore la santé publique, réduit les émissions liées aux trajets motorisés courts et renforce l'accessibilité urbaine [1]. Pourtant, dans de nombreuses villes, sa pratique quotidienne reste limitée, notamment parce que beaucoup d'usagers potentiels le jugent dangereux ou inconfortable en trafic mixte [2]. La littérature montre de façon convergente que le degré de séparation entre cyclistes et trafic motorisé est un déterminant central de la sécurité, perçue comme objective, et du confort : plus la séparation physique est forte, plus la propension à se déplacer à vélo augmente [3, 4]. Dans ce contexte, les collectivités doivent décider où investir pour maximiser les gains de sécurité et de confort, tout en respectant des budgets limités et des contraintes de mise en oeuvre.

Cette thèse propose une suite de modèles d'aide à la décision pour améliorer la sécurité et le confort cyclables à l'échelle d'un réseau, sous contrainte budgétaire. Le choix méthodologique central est de représenter l'infrastructure cyclable par un graphe orienté pondéré construit à partir de données OpenStreetMap : les arcs représentent des segments et leurs poids codent la qualité de séparation (DoS). Les investissements sont modélisés comme des améliorations discrètes du DoS sur un sous ensemble d'arcs. Cette représentation permet de progresser par étapes : une planification fondée sur la structure du réseau, puis une planification sensible à la demande via la redistribution des flux, ensuite une extension multimodale qui explicite les rétroactions sur le trafic automobile lorsque l'espace routier est réalloué, puis une intégration du report modal et d'objectifs de justice des transports, et enfin un cadre de comparaison inter villes fondé sur une réduction de réseau afin de limiter la sensibilité à la taille et à la résolution cartographique.

Une idée directrice traverse l'ensemble des chapitres : la sécurité cyclable n'est pas seulement une propriété locale d'un tronçon, mais un phénomène de réseau. Elle dépend de la continuité des itinéraires protégés, de la diversité des itinéraires disponibles, de la distribution spatiale de la demande, des interactions entre modes et des objectifs distributifs. L'enjeu n'est donc pas uniquement d'identifier des segments à risque, mais de comprendre quelle structure rend possibles des déplacements sûrs et continus, quels investissements bénéficient réellement à une part importante des cyclistes, et comment préserver l'équité tout en gérant les effets sur les autres modes.

Chapitre 3: Conception de réseau fondée sur la topologie

Le premier chapitre pose les bases d'un cadre de planification fondé sur la structure du réseau cyclable. À partir de données OpenStreetMap, l'infrastructure est modélisée comme un graphe orienté pondéré, et les arcs sont classés selon un schéma DoS à six niveaux. Ces niveaux sont ensuite alignés avec CycleRAP afin d'obtenir une pondération cohérente orientée sécurité et confort [5]. Le chapitre traite un point méthodologique essentiel pour l'optimisation : si l'on modélise un investissement comme une augmentation du DoS, les indicateurs utilisés pour

guider la décision doivent évoluer de manière cohérente, idéalement sans se dégrader lorsque l'infrastructure devient plus sûre. Sans cette propriété, une procédure d'optimisation peut être biaisée et recommander des solutions difficiles à interpréter.

Le chapitre propose une analyse systématique de métriques de théorie des graphes, souvent utilisées pour décrire la robustesse et la structure des réseaux [6], et étudie leur comportement sous des améliorations aléatoires. Cette étude met en évidence un ensemble d'indicateurs particulièrement adaptés au contexte cyclable : la force de noeud pondérée, une centralité de proximité définie au niveau des arcs via un graphe ligne, et des mesures de connectivité telles que la plus grande composante connexe d'infrastructures fortement séparées. Ensemble, ces indicateurs décrivent si les segments sûrs forment un réseau utilisable, si des corridors structurants existent, et si l'on observe une continuité plutôt qu'une juxtaposition de fragments.

Le chapitre relie ensuite ces indicateurs à la sécurité et au confort perçus en les confrontant à des données d'enquête (Baromètre Vélo), afin de construire un modèle explicite reliant structure de réseau et perception. Sur cette base, un problème d'optimisation sous contrainte budgétaire est formulé : sélectionner un ensemble d'améliorations maximisant la sécurité et le confort estimés. La méthode est comparée à plusieurs solveurs heuristiques, notamment des algorithmes génétiques et l'optimisation bayésienne, et, lorsque possible, à une référence exhaustive. Sur 14 villes françaises, les résultats indiquent des améliorations systématiques et un compromis favorable entre qualité de solution et coût de calcul. Un enseignement important est que relier des segments déjà protégés peut produire des gains significatifs à coût modéré, car cela transforme des aménagements isolés en réseau réellement praticable.

Chapitre 4: Conception de réseau avec redistribution des flux

Le deuxième chapitre montre que la topologie ne suffit pas toujours pour prioriser les investissements. Un lien peut être central d'un point de vue structurel tout en étant peu utilisé ; l'améliorer peut alors avoir un impact limité. L'objectif devient donc d'intégrer la demande et les choix d'itinéraires pour passer d'une logique de structure à une logique d'usage effectif.

Le chapitre introduit une affectation cycliste de type all or nothing : chaque paire origine destination est affectée au plus court chemin selon un coût généralisé adapté au vélo, qui combine longueur, pente, changements d'angle et un facteur lié aux conditions de trafic. On en déduit des flux cyclistes sur les arcs et une mesure d'exposition qui pondère le risque par l'usage. Un segment très exposé mais rarement emprunté ne pèse pas autant qu'un segment moins exposé mais traversé par de nombreux cyclistes. Le problème prend une structure bi niveau : les améliorations modifient les coûts, les itinéraires les plus courts changent, les flux se redistribuent et l'exposition évolue.

Sur Grenoble, les résultats montrent que l'intégration des flux change les corridors prioritaires par rapport aux stratégies purement topologiques. À budget fixé, l'optimisation cible des arcs combinant faible DoS, forte exposition et forte utilisation prédite. Les plans obtenus réduisent l'exposition agrégée et réorientent une part significative des cyclistes vers des itinéraires mieux séparés. La comparaison algorithmique indique que l'algorithme génétique

modifié identifie des portefeuilles d'améliorations très efficaces, notamment en exploitant des améliorations intermédiaires de DoS offrant de bons gains par unité de coût. À l'inverse, une stratégie centrée sur l'extension de la composante la plus sûre peut améliorer la connectivité des meilleurs niveaux mais rester moins alignée avec la demande, et donc moins efficace pour réduire l'exposition là où les cyclistes circulent réellement.

Chapitre 5: Conception multimodale avec redistribution des flux

Le troisième chapitre traite une contrainte centrale en pratique : les aménagements cyclables sont souvent réalisés par réallocation d'espace, ce qui peut réduire la capacité automobile et augmenter les temps de parcours. Les décideurs doivent donc arbitrer entre gains de sécurité à vélo et impacts sur la circulation motorisée, surtout dans des environnements urbains contraints.

Le chapitre propose un cadre bi niveau multimodal. Au niveau supérieur, le planificateur choisit des améliorations de DoS afin de réduire l'exposition des cyclistes tout en limitant la dégradation des temps de trajet automobiles liée aux changements de capacité. Au niveau inférieur, les cyclistes choisissent leurs itinéraires selon le coût généralisé, tandis que les temps automobiles sont évalués avec un modèle de trafic cohérent avec la capacité modifiée. Le cadre permet ainsi d'évaluer chaque portefeuille d'investissements du point de vue des bénéfices cyclables et de la faisabilité opérationnelle.

Sur Grenoble, les résultats suggèrent que des réductions importantes d'exposition cycliste sont souvent atteignables avec des augmentations limitées du temps total automobile, à condition de cibler des corridors à forte demande cycliste et d'éviter les axes automobiles déjà critiques. Le modèle rend explicites ces compromis et met en évidence qu'il existe des options proches en bénéfice cyclable mais très différentes en coût automobile. Il fournit donc un appui pour sélectionner des plans robustes, au sens où ils restent acceptables pour le fonctionnement global du système de transport.

Chapitre 6: Conception multimodale pour la justice des transports

Le quatrième chapitre étend le cadre précédent dans deux directions. D'une part, il intègre le report modal : améliorer la sécurité et le confort peut inciter une fraction d'automobilistes à passer au vélo, ce qui modifie la demande et peut atténuer certains effets de congestion. D'autre part, il introduit des objectifs de justice des transports, en considérant que la planification doit aussi garantir des standards minimaux de sécurité et d'accessibilité, et pas seulement améliorer des moyennes.

Le chapitre s'appuie sur la littérature en justice des transports [7] et formalise deux perspectives. La perspective utilitariste vise à améliorer la performance globale, par exemple en réduisant l'exposition totale tout en limitant l'augmentation des temps automobiles. La perspective suffisantariste cherche à réduire la part d'usagers sous un seuil acceptable, par ex-

emple des cyclistes qui restent contraints à des itinéraires trop exposés ou des automobilistes qui subissent des temps excessifs. Pour rester calculable à l'échelle urbaine, un modèle surrogate est utilisé afin d'approximer les flux automobiles proches d'un équilibre utilisateur, tout en capturant les rétroactions de capacité et de demande.

Les expériences sur Grenoble soulignent trois points. Premièrement, le report modal modifie l'évaluation des impacts automobiles : des réductions de capacité peuvent devenir plus acceptables si une partie des trajets bascule vers le vélo. Deuxièmement, la différence entre utilitarisme et suffisantarisme est nette : de bonnes performances globales peuvent coexister avec des zones ou des groupes encore sous standard. Troisièmement, l'approche suffisantariste produit des solutions plus équilibrées, en réduisant la part d'usagers sous seuil tout en conservant des résultats compétitifs au niveau agrégé. Le chapitre montre ainsi qu'une planification plus équitable peut rester compatible avec des objectifs d'efficacité, à condition d'intégrer explicitement la demande induite et des contraintes minimales.

Chapitre 7: Regrouper et comparer la sécurité des réseaux cyclables

Le cinquième chapitre traite un problème méthodologique majeur pour les comparaisons inter villes. Les réseaux diffèrent en taille, en couverture et en résolution cartographique ; comparer les graphes bruts peut donc confondre différences de représentation et différences réelles de structure de sécurité. Il est alors nécessaire de construire des représentations comparables, qui conservent l'information pertinente tout en réduisant la sensibilité à la résolution.

Le chapitre propose une réduction de réseau par clustering contraint et orienté sécurité. Un grand réseau est compressé en un super réseau de taille fixée, avec le même nombre de super noeuds qu'un réseau de référence. La méthode s'appuie sur un plongement spectral et un k means modifié, avec des contraintes visant à préserver DoS, distances, pentes, facteurs de trafic, flux et exposition. Chaque cluster devient un super noeud, et les super arcs agrègent les connexions inter zones, de manière à conserver l'organisation des corridors protégés et leurs ruptures.

Sur ces réseaux réduits, le chapitre introduit des mesures de similarité fondées sur l'exposition et la structure, notamment une comparaison de type transport optimal (Gromov Wasserstein) et une distance entre noyaux d'adjacence pondérés par l'exposition. Ces outils ne dépendent pas du renommage des noeuds et restent sensibles aux poids de sécurité, ce qui permet des comparaisons plus justes entre villes. Les résultats montrent que le clustering contraint préserve mieux flux et exposition qu'un clustering non contraint, et que les distances obtenues sont stables et interprétables. Ce cadre fournit une base pour comparer des villes, suivre des évolutions dans le temps et identifier des villes comparables pour l'échange de bonnes pratiques.

Conclusions: Synthèse et apports

Cette thèse propose une chaîne méthodologique cohérente pour planifier la sécurité cyclable sous contraintes réalistes. Les contributions se regroupent en quatre ensembles. Premièrement, elle formalise la sécurité à vélo comme un phénomène de réseau et identifie des indicateurs adaptés à une logique d'amélioration progressive, avec un comportement cohérent sous upgrades. Deuxièmement, elle relie infrastructure et usage en intégrant une affectation cycliste, des flux et une mesure d'exposition afin de prioriser des investissements là où ils bénéficient réellement aux cyclistes. Troisièmement, elle étend l'évaluation à la multimodalité et rend explicites les compromis entre sécurité cyclable et performance automobile, tout en intégrant le report modal. Quatrièmement, elle propose un cadre de réduction et de comparaison inter villes qui limite l'influence de la taille et de la résolution cartographique sur les conclusions.

La conclusion générale est que l'amélioration efficace de la sécurité cyclable repose sur la construction d'une structure de réseau continue, attractive et équitable, et pas seulement sur des interventions locales. En articulant topologie, flux, multimodalité, justice et comparaison inter villes, la thèse fournit une base robuste pour l'aide à la décision.

Perspectives de recherche et d'extension

Les travaux présentés ouvrent plusieurs prolongements naturels pour renforcer à la fois la crédibilité empirique des indicateurs, la capacité de passage à l'échelle et la pertinence pour l'action publique. Un premier enjeu consiste à consolider les bases empiriques de la variable centrale de décision, le DoS, en améliorant et en validant la calibration DoS–CSCS lorsque des données locales existent : l'objectif est de vérifier que des valeurs de DoS plus élevées et une exposition prédite plus faible se traduisent effectivement par moins d'accidents, moins de situations de conflit ou une meilleure perception de sécurité, et d'identifier les cas où les informations OpenStreetMap sont insuffisantes, en particulier aux intersections ou sur de courtes zones de mélange, afin de proposer des ajustements ciblés. Dans la même logique, un second axe important est la validation directe de l'indicateur d'exposition, qui joue un rôle structurant dans les chapitres orientés flux : confronter l'exposition estimée à des sources indépendantes, comme des mesures de perception du danger, des données d'accidents et de blessures ou des quasi accidents, permettrait de renforcer l'interprétation des objectifs d'optimisation et de repérer d'éventuels biais liés à certaines typologies d'axes ou de contextes urbains.

Sur le plan comportemental, une évolution attendue est de dépasser l'affectation all or nothing, efficace mais restrictive, afin de mieux représenter la diversité des choix d'itinéraires et l'hétérogénéité des préférences. Des modèles stochastiques intégrant plusieurs classes d'usagers, avec des sensibilités différenciées au DoS et au trafic [8], permettraient d'évaluer les investissements non seulement via des déplacements de plus courts chemins, mais aussi via la diversification des routes et l'évolution de la population prête à se déplacer à vélo au fur et à mesure que le réseau s'améliore. Cette amélioration du réalisme peut aussi s'appuyer sur une prise en compte plus explicite de la dimension temporelle : la demande varie selon l'heure, la saison et la météo, et la congestion est intrinsèquement dynamique. Évaluer les

plans d'investissement sur plusieurs conditions d'exploitation et sur des scénarios contrastés renforcerait la robustesse des recommandations, et enrichirait les comparaisons inter villes en introduisant une notion de similarité non seulement moyenne, mais aussi en robustesse.

Par ailleurs, l'extension multimodale peut être approfondie en intégrant la marche et le transport public, ainsi que des contraintes d'espace telles que stationnement et livraisons. Une telle évolution ferait progressivement basculer l'évaluation d'une logique centrée sur la qualité d'infrastructure vers une logique d'accessibilité aux emplois et services, et ouvrirait la voie à des objectifs de justice plus directement opérationnels, en cohérence avec des travaux récents [9, 10]. Dans le même temps, des modèles plus fins des effets de capacité, par exemple liés aux feux, aux usages de voie ou à des contraintes locales, amélioreraient l'analyse des compromis entre gains cyclables et impacts automobiles. Enfin, pour traiter des réseaux de grande taille et des contraintes plus riches, la réduction par clustering suggère une stratégie de passage à l'échelle particulièrement prometteuse : résoudre d'abord un problème d'investissement sur un super réseau, puis remonter les décisions vers le réseau complet, et, à terme, mettre en place une boucle multi résolution alternant décisions grossières et raffinements fins pour accélérer l'exploration de l'espace de solutions sans perdre les effets locaux.

Ces prolongements méthodologiques gagnent toutefois à être articulés avec la planification réelle. Les décisions d'aménagement dépendent de gouvernance, d'acceptabilité et de phasage, et les recommandations doivent pouvoir intégrer des contraintes de faisabilité, produire des analyses distributives détaillées et s'accompagner d'analyses d'incertitude. Dans cette perspective, la comparaison inter villes peut être formulée comme un outil d'aide à la décision, au service de l'identification de villes comparables, d'interventions transférables et d'indicateurs de suivi. En synthèse, la thèse propose un pipeline allant d'indicateurs topologiques adaptés aux améliorations vers une optimisation fondée sur les flux et l'exposition, puis une extension multimodale orientée justice, et enfin une comparaison inter villes rendue possible par une réduction préservant l'information de sécurité. L'étape la plus immédiate consiste à exploiter le clustering pour accélérer l'optimisation tout en conservant la fidélité sécurité ; à plus long terme, la validation empirique, l'enrichissement comportemental, la prise en compte du temps et l'intégration de contraintes de mise en oeuvre rapprocheront ce cadre d'un outil opérationnel pour des réseaux cyclables plus sûrs, plus efficaces et plus équitables.

Acknowledgments

First, I would like to apologize for the length of this section. I had hoped to keep it shorter, but everyone mentioned here has, in one way or another, shaped my life during these years. I also apologize to anyone I may have unintentionally omitted.

I would like to express my deepest gratitude to my supervisors, Prof. Carlos Canudas de Wit and Dr. Giovanni De Nunzio, for giving me the opportunity to work under your guidance. Thank you for your constant encouragement and advice during the most difficult moments, and for the countless times you read my drafts and provided clear, precise, and constructive feedback, even during weekends and holidays. I am especially grateful for your ability to turn my rough and sometimes messy ideas into structured contributions, and for always listening to me with care and patience, even when it was not easy. I hope that one day I can become, even in a small way, the kind of researcher you are.

To the professors and researchers of the DANCE team (Giacomo, Hassen, Federica, Alain, and Gustave), thank you for your kindness and support. I also thank the members of my jury, Sophie Achard, Jeppe Rich, Luca Consolini, Alain Kibangou, Thomas Buhler, and Mauro Salazar for accepting to take part in my defence and for taking the time to read my thesis.

I am especially grateful to Prof. Mauro Salazar for welcoming me in Eindhoven and for giving me the opportunity to work with the MOVEMENT group. I only have excellent memories of that visit and of the generosity of your team. Thank you as well for your ideas to improve our models and for introducing us to the rich and important topic of transport justice.

To my colleagues and friends at Inria, Léo, Baptiste, Rémi, and Jean, thank you for your kindness and for your help during the time we overlapped. Martín, thank you for your advice and for helping me understand the french system at the beginning of the Ph.D. Mladen, thank you for your camaraderie from the very first day we met, and for continuing to show your kindness at the conferences where we crossed paths. Thank you also for always eating the most whenever I cooked. Tommaso, thank you for being so kind, and for helping me with phone calls in French when I struggled with administrative matters. I miss our conversations about anime, and I also thank you for introducing me to Giovanni, Pietro, and Alex, it does not surprise me that your friends are as cool and kind as you are.

Yohan, thank you for your kindness and for offering your help on several occasions; you are a sharp and genuinely interesting person. Sébastien, thank you for speaking with me so naturally, despite how intelligent you are; I am sure you will become an excellent professor. Guillaume, I am grateful that you are the french person with whom I have spent the most time. I once told a french professor how selflessly you helped me and how you regularly checked in, and she told me that it is rare and that I should take care of this friendship. I still regret declining your very nice invitation to visit your home. One thing I would have loved is to have

been a better friend to you. I am sure you will do very well in Lyon; I have always admired your quick mind. I wish you and your partner, Aude, all the very best. Yohan, Sébastien, and Guillaume, thank you for your patience when speaking french with me, and I am sorry for the effort it sometimes took to understand what I was trying to say, I still have a long way to go.

To my Latin friends, I want to repeat something I have said many times and truly mean: Somehow, your friendship makes me feel closer to home. Estela, thank you for being so kind and for introducing me to your friends; I am sure you will successfully finish your Ph.D. and become a great researcher in Peru. Manuela, thank you for your kindness, especially for taking me on a bike trip to see a lake when I was struggling; it is a memory I will always cherish. And your cooking is amazing. Ignacio, thank you for being so friendly and for listening to me; thank you for coming with me to the hospital when I needed tests, and for inviting me to bouldering. Talking with you helped me realize that I would like to expand my life and spend more time outside of my usual habits. Vicente, I truly mean what I once told you: you shine when you sing in your choir. Please never stop, every time I have heard you, it made me feel full. Thank you for teaching me the basics of salsa and cumbia, and for welcoming us into your home. Your cooking is better than any mexican restaurant I have tried outside Mexico, and I always enjoyed talking and eating with you. Eduardo, thank you for being my friend. Despite being so prudent, funny, cultured, and thoughtful, you still chose to spend time with someone as ignorant as me. Thank you for inviting me to meals, museums, and film presentations; I have learned so much from you, especially to be more prudent, empathetic, patient, and to live more calmly.

I also thank the local church in Grenoble for welcoming me as a member, especially Salvo. I apologize for being inconsistent during the Ph.D., but I often needed Sundays to work from home.

A la iglesia Amisadaí el Chacón, Pastores, Suriel, Bety, Gabriel, Julieta, Juanita, Abel, Tomás Gonzales y Juan Carlos, gracias por ser un apoyo para mí y mi familia.

A mis compañeros y amigos de México, Eric, Diana, Are, Luis, Alexis, Arturo, Héctor y Katia: gracias por estar siempre presentes, incluso a la distancia. Eric, gracias por ser un verdadero amigo durante tanto tiempo. Katia, gracias por escuchar con tanta atención mis historias, creo firmemente que tu amabilidad es tu don más grande y contigo una conversación breve se siente como volver a casa. Tal vez suene exagerado, pero muchas veces un par de mensajes suyos fue suficiente para cambiarme el ánimo y alegrarme el día.

A mi papá (Amado Adrián Campero Villa), gracias por haber trabajado en cosas tan físicamente demandantes para que nunca faltara comida en la casa. Yo nunca comprendí realmente lo difícil que era hasta el día que comencé a acompañarte a trabajar. Y a mi mamá (Blanca Jurado Zaragoza), el cariño que nos has dado es algo que nunca podré pagar. Gracias por cuidarnos y por hacer que la casa esté siempre llena de música y de cantos gracias a ti. Los amo e intento hacerlos sentir orgullosos.

Gali (Galilea Campero Jurado), tú sin duda eres el miembro más importante y valioso de nosotros tres. Hablar contigo y ver a la Chuby y Kayn siempre me robaba sonrisas sinceras. Gracias por ser tan leal y responsable, y por cuidar de Papá, Mamá y a los gordos en nuestra ausencia. Estamos realmente orgullosos de ti y te amamos.

Abuelo (Guillermo Campero Ramírez), usted ya no se encuentra físicamente, pero gracias por todas las cosas que hizo por nosotros. Mientras viva, jamás olvidaré su amabilidad y gentileza, e intentaremos cuidar a mi Abuela y a mi tía.

Bebo NVIDIA (Israel Campero Jurado), siempre he pensado que tenerte como hermano es la mayor bendición que he tenido en esta vida. Tú no sólo eres mi mejor amigo, sino también el mejor amigo que pienso que alguien puede tener. Tu esfuerzo, hambre de salir adelante y preocupación por tu familia son cosas a las que aspiro llegar a ser. Tú siempre has estado conmigo, en mis peores momentos, y nunca me abandonaste; la frase "sabía que volverías por mí" resume perfectamente la lealtad que te define. Gracias a ti, siento que aún puedo convertirme en la persona que yo quería ser cuando era un niño.

Finalmente, gracias, Dios, por la oportunidad inmerecida de haber vivido esta aventura y de conocer a personas tan asombrosas. Aun con mis errores, y tan lejos de casa, nunca dejé de sentir tu misericordia.

"Si tomare las alas del alba y habitare en el extremo del mar, aun allí me guiará tu mano, y me asirá tu diestra."— Salmo 139:9–10

Contents

Résumé **i**

List of acronyms **xxiii**

List of symbols **xxv**

1 Introduction **1**

1.1 Cycling for Healthy and Accessible Urban Mobility 1

1.2 Benefits of Cycling 2

1.3 Barriers to Cycling 5

1.4 Literature Review 7

1.5 Thesis Motivation and Contributions 19

1.6 Publications 21

1.7 Thesis outline 22

2 Preliminaries **27**

2.1 Cycling Network Definition and Classification 27

2.2 Cyclist Generalized Distance 29

2.3 Flow Assignment in Bicycle and Car Layers 30

2.4 Safety and Efficiency Metrics 31

2.5 Objective Functions 31

3 Topological graph-based network design **35**

3.1 Introduction 35

3.2 Centrality Measures for Safety and Comfort 39

3.3 Model Design 49

3.4 Solution Algorithms . . . 57
3.5 Benchmark Instances and Case Studies . . . 61
3.6 Concluding Remarks . . . 68

4 Network Design with Flow Redistribution 71

4.1 Introduction . . . 71
4.2 Model Design . . . 73
4.3 Case Study . . . 78
4.4 Concluding Remarks . . . 82

5 Network Design with Flow Redistribution in a Multi-Modal Framework 85

5.1 Introduction . . . 85
5.2 Model Design . . . 87
5.3 Case Study . . . 93
5.4 Concluding Remarks . . . 102

6 Network Design with Flow Redistribution in a Multi-Modal Network for Transport Justice 103

6.1 Introduction . . . 103
6.2 Model Design . . . 107
6.3 Case Study . . . 118
6.4 Concluding Remarks . . . 128

7 Clustering and Comparing Safety in Cycling Networks 131

7.1 Introduction . . . 131
7.2 Model Design . . . 134
7.3 Clustering (Doc-CS) . . . 138
7.4 Network Comparison . . . 143
7.5 Benchmark Instances and Case Study . . . 144

7.6 Concluding Remarks 149

8 Conclusions and Future Research **151**

A Complementary proofs **157**

Bibliography **178**

List of Figures

2.1 At the top: Cycling network of Grenoble, France, with the DoS ranging from 1 to 6, indicating increasing separation from motorized traffic. At the bottom: Road network for cars, including links shared with bicycles. 29

3.1 Perceived cycling safety and comfort data in Grenoble, France. 38

3.2 CycleRAP demonstrator tool for scoring cycling facilities. 41

3.3 Comparison of DoS-weighted centrality measures: a, strength centrality. b, k-core. c, closeness centrality. d, betweenness centrality e, eigenvector centrality. f, accessibility. The size and darkness of each node are proportional to its centrality measure and the color of the arcs is according to its DoS: 6: ●, 5: ●, 4: ●, 3: ●, 2: ●, 1: ●. 44

3.4 Monotonicity behavior of weighted graph metrics under 1000 random DoS upgrades. "Up" indicates that the graph metric increased after the upgrade, "Down" that it decreased, and "Zero" that it remained unchanged. 46

3.5 Flowchart of the model training and testing process. 52

3.6 Predicted values versus the actual values of perceived safety (right) and comfort (left). The average error is 2.23 % and 3.15 % for Auto-ML Ensemble and NNLS, respectively, for safety, and 5.99 % and 7.36 % for Auto-ML Ensemble and NNLS, respectively, for comfort. 55

3.7 (a) Original configuration, (b) optimal solution configuration, (c) MLCC algorithm configuration, (d) GA configuration, (e) GA-U configuration and (f) Bayesian algorithm configuration: These configurations aim to enhance perceived safety with $i = 7$, corresponding to $\alpha_7 \approx 0.35\%$ of the total budget required for a full network upgrade ($B_{\max}$ =22,010€). 62

3.8 (a) Results of solving Problem 3.2 using the optimal, MLCC, GA, GA-U, and Bayesian algorithms on the benchmark network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\text{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\text{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) Performance evolution of the GA and GA-U algorithms over generations for safety with (c) providing the legend. $B_{\max} = 22,010$ €. 63

3.9 Calculation time for the benchmark example 63

3.10 Sioux Falls cycling network 65

3.11 (a) Results of solving Problem 3.2 using the MLCC, GA, GA-U, and Bayesian algorithms on the Sioux Falls cycling network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) show the performance evolution of the GA and GA-U algorithms over generations for safety with (c) providing the legend. And $B_{\max} = 194{,}848€$. 66

3.12 Calculation time for the Sioux Falls cycling network. 66

3.13 Cycling network of Grenoble, France, after being upgraded using MLCC with $\alpha_7 \approx 0.35$. 67

3.14 (a) Results of solving Problem 3.2 using the MLCC, GA, GA-U, and Bayesian algorithms on the Grenoble cycling network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) Performance evolution of the GA and GA-U algorithms over generations for safety, with (c) providing the legend. 68

3.15 Calculation time for the Grenoble cycling network. 68

3.16 Results of the optimization problem (MLCC algorithm) for safety (top) and comfort (bottom) for logarithmically spaced values of α. $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. 69

4.1 Layout of Grenoble city, origin/destinations nodes in regions $\mathcal{R}$. 75

4.2 On the top is the initial normalized traffic exposure and reduced exposure after optimization using the GA-U and MLCC algorithms for Problem 4.1 with a budget constraint of 50%, i.e., $\alpha = \frac{1}{2}$. And on the bottom the convergence analysis of GA-U for the same budget. 79

4.3 The upper row shows the DoS configuration, and the lower row shows the flow mapping. The optimization algorithms operate with 50% of the total budget needed for a complete network upgrade. 81

4.4 Distribution of bike-lane length by DoS (normalized) 81

4.5 Global measure of perceived safety and comfort. 82

5.1 Multi-modal transportation network of Grenoble, France. *Left:* Cycling infrastructure classified by degree of separation (DoS) of bike lanes, ranging from 1 (least separated) to 6 (fully separated). *Right:* Road network for cars, including links shared with bicycles. 87

5.2 BPR function. 88

5.3 The creation/upgrading of a bike lane is modeled as reducing the capacity of the affected car link by half, resulting in an increased travel time $\uparrow t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}})$. . . . 90

5.4 Original pre-upgrade UE flows (normalized) obtained by solving Problem 5.1. . 94

5.5 Normalized travel time versus cyclist exposure. The top plot displays the Pareto fronts resulting from the trade-off Problem 5.2 across different values of $\gamma \in \Lambda$ and $\alpha \in \Psi$. The medium plot presents a zoomed-in view focusing on the optimal γ-constraint solutions. The bottom figure present the convergence analysis of the GA-U algorithm for solving Problem 5.2 on the optimal points $\gamma^\star_\alpha$. 97

5.6 Normalized cyclist exposure and travel time for each $\gamma \in \Lambda$ and $\alpha \in \Psi$. In red color are the $\gamma^\star_{\frac{1}{3}}, \gamma^\star_{\frac{2}{5}}$ and $\gamma^\star_{\frac{1}{2}}$ from top to bottom, respectively. 98

5.7 Comparison of car travel times at the optimal trade-off points $\gamma^\star_\alpha$, with post-upgrade UE traffic assignment after bike lane upgrades at these locations. . . . 99

5.8 GA-U upgrades and post-upgrade user-equilibrium (UE) assignments for three budgets. 100

5.9 MLCC upgrades and post-upgrade user-equilibrium (UE) assignments for three budgets. 101

6.1 The creation/upgrading of a bike lane is modeled as reducing the capacity of the affected car link, resulting in an increased travel time $\uparrow t^{\mathrm{c}}_a(\phi^{\mathrm{c}}_a)$. 108

6.2 Historical cycling network of Grenoble, France. *Left:* year 2015. *Right:* year 2025. The average cycling distance $\overline{D}^{\mathrm{b}}[k]$ decreases as the network expands and improves. 109

6.3 The logit model plot to predict $P^{\mathrm{b}}[k]$ using $\bar{D}^{\mathrm{b}}$. 110

6.4 On the top we present the epochs evolution with the 912 samples created for solving Problem 5.1. On the middle how the predicted flows approaches the true flow found by the UE, and on the. And on the bottom we present the residual analysis. 119

6.5 Normalized travel time versus cyclist exposure. The top panel shows the Pareto fronts obtained from the trade-off Problem 6.1 under varying values of $\gamma \in \Lambda$ and $B \approx 0.15 B_{\max}$. The medium panel presents the corresponding results for the sufficientarian formulation Problem 6.2. And the bottom row present the convergence analysis for the utilitarian optimal and sufficientarian for $\lambda = \frac{2}{5}$. . . 121

6.6 Normalized cyclist exposure, travel time, safety path insufficiency, cyclist and drivers trip insufficiency trade-off for $\lambda = \frac{2}{5}$. 122

6.7 Path safety versus the number of cyclists above and below the threshold $\lambda = \frac{2}{5}$ under a budget of $B \approx 0.15 \cdot B_{\max}$. The dotted blue and red lines indicate the relative path-safety threshold s_{thr} and the average safety score of the entire network $S^{\mathrm{b}}_{\mathcal{G}_{\mathrm{avg}}}$, respectively. 122

6.8 Drivers (left y-axis) and cyclists (right y-axis) trip insufficiency versus the number of individuals above and below the threshold $\lambda = \frac{2}{5}$ under a budget of $B^*_\lambda \approx 0.15 \cdot B_{\max}$ for the different optimization algorithms. The dotted red, blue, and green lines indicate the trip-insufficiency threshold T_{thr}, the average bicycle trip insufficiency $T^{\text{b}}_{\text{avg}}$, and the average car trip insufficiency $T^{\text{c}}_{\text{avg}}$ of the entire network, respectively. 123

6.9 Normalized path-safety insufficiency, bike path-time insufficiency and car path-time insufficiency for $\lambda = \frac{2}{5}$. 124

6.10 DoS configurations (top row) and regional safety path insufficiency (second row) for the original network and the MLCC, GA-U, and GA-S upgrade plans. All methods use $\lambda = \frac{2}{5}$ and budget $B \approx 0.15\, B_{\max}$. 125

6.11 Regional bicycle travel time insufficiency (top row) and car travel time insufficiency (bottom row) for the original network and the MLCC, GA-U, and GA-S upgrade plans. All methods use $\lambda = \frac{2}{5}$ and budget $B \approx 0.15\, B_{\max}$. 126

6.12 Safety insufficiency comparison of upper level $\mathbb{U}_{\text{thr}}$ of Problem 6.2, and the color legend indicates the λ threshold. 127

6.13 Normalized cyclist exposure. 128

6.14 Total system travel time ($\mathcal{TTT}^{\text{tot}}$). 128

6.15 Grenoble configuration paths under different optimizations algorithms. 129

7.1 DoS-CS Algorithm for toy example where $a = \{w_a, d_a(\text{m}), \phi_a, \mathcal{E}_a\}$, $\tilde{a}^{\text{inter}} = \left\{w^{\text{inter}}_{\tilde{a}}, d^{\text{inter}}_{\tilde{a}}(\text{m}), \phi^{\text{inter}}_{\tilde{a}}, \mathcal{E}^{\text{inter}}_{\tilde{a}}\right\}$, and $C^{\text{intra}} = \left\{d^{\text{intra}}_{C}, d^{\text{intra}}_{C}(\text{m}), \phi^{\text{inter}}_{C}, \mathcal{E}^{\text{intra}}_{C}\right\}$ 139

7.2 Academic example for clustering process 145

7.3 Evolution of cyclist exposure and cyclist flow through the DoS-CS algorithm (top: Lyon; bottom: Strasbourg). 146

7.4 Clustering process with the DoS-CS algorithm. The first column shows the original networks, the second the k-means clustering, and the third the DoS-CS clustering (top: Lyon, bottom: Strasbourg). 148

A.1 Flowchart of the MLCC algorithm. 161

List of Tables

1.1 DoS levels derived from OSM tags and their alignment with the CycleRAP Safety and Comfort Score (CSCS), illustrated with Grenoble traffic context (AADT and mean speed near cycling links). Higher DoS indicates stronger separation from motor traffic and corresponds to higher CSCS values. 11

1.2 Unified summary of input variables used by bicycle facility assessment methods. A checkmark indicates explicit use of the variable in the method (columns 1–12 correspond to roadway/traffic and context variables; columns 13–23 cover additional operational and service variables). 12

1.3 Representative optimization/prioritization approaches for bike network design and expansion. 16

3.1 The table displays the degree of separation (DoS) with its CycleRAP Safety and Comfort Score (CSCS), cycling infrastructure category, and GTLVille data on AADT and mean speed for roads near cycling paths in Grenoble. 40

3.2 Values of $s_{\mathcal{G}^{\mathrm{b}}}, c_{\mathcal{G}^{\mathrm{b}}}$ and $l_{\mathcal{G}^{\mathrm{b}}}$ for fourteen French municipalities. Safety (y_s) and comfort (y_c) values are scaled to the interval $[0, 1]$. 51

3.3 Regression results for solving Eq. (3.34) for safety and comfort. 52

3.4 Comparison of the most relevant models using AICc and LOOCV-RMSE. Lower values are better for both criteria. 54

3.5 The cities, the number of nodes, arcs, total amount of meters and the budget required to upgrade all arcs to the highest possible DoS. 67

4.1 Notation summary of the Chapter 4. 74

4.2 Abbreviated notation for variables in cycling network model, writing $(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ next to a variable indicates that it is a function of the candidate configuration for bike lane improvements. 76

4.3 Slope factors for cycling distance calculation 78

5.1 Abbreviated notation for variables in cycling network model, writing $(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ next to a variable indicates that it is a function of the candidate configuration for bike lane improvements. 92

6.1 Notation used in the utilitarian optimization formulation. 113

6.2 Abbreviated notation for variables in trade-off network model. Notation for optimization variables depending on the upgrade configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$ 117

List of acronyms

Acronym	Meaning
DoS	Degree of Separation
CSCS	CycleRAP Safety and Comfort Score
FS	Fully Safe network (all DoS = 6)
GTLVille	Grenoble Traffic Lab Ville Platform
CC	Connected Component
SDV	Synthetic Data Vault
MLCC	Maximise Largest Connected Component (algorithm)
NNLS	Non-Negative Least Squares
OSM	OpenStreetMap
AADT	Annual Average Daily Traffic
AoN	All-or-Nothing traffic assignment
UE	User equilibrium
BPR	Bureau of public roads function
GCN	Graph Convolutional Network
DoS CS	Constrained spectral clustering (algorithm)

List of symbols

$\mathcal{G}^{\mathrm{b}} = \big(\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}}\big)$	Cycling network as a weighted directed graph (bike layer).
$\mathcal{V}^{\mathrm{b}}$	Set of cycling nodes (intersections / road endpoints).
$\mathcal{A}^{\mathrm{b}}$	Set of directed cycling arcs (bike-lane segments).
$\mathcal{W}^{\mathrm{b}}$	Set of cycling arc weights (e.g., DoS / CSCS-derived values).
$\mathcal{G}^{\mathrm{c}} = \big(\mathcal{V}^{\mathrm{c}}, \mathcal{A}^{\mathrm{c}}, \mathcal{W}^{\mathrm{c}}\big)$	Car network as a weighted directed graph (car layer).
$\mathcal{V}^{\mathrm{c}}$	Set of car nodes (intersections / road endpoints).
$\mathcal{A}^{\mathrm{c}}$	Set of directed car arcs (road segments for cars).
$\mathcal{W}^{\mathrm{c}}$	Set of car arc weights (e.g., travel-time / capacity-related attributes).
$a = (u, v)$	Directed arc from node u to node v (in the relevant layer).
d_a^{b}, d_a^{c}	Geographical length of arc a in the bike / car layer.
w_a^{b}	Degree of Separation (DoS) level of cycling arc a (higher is safer).
$h : \{1, \ldots, 6\} \to \mathbb{R}_{>0}$	Mapping from DoS level to CSCS weight.
$h(w_a^{\mathrm{b}})$	CSCS value associated with cycling arc a.
$\mathcal{N}_v^{\mathrm{b}}$	Set of cycling arcs incident to a cycling node $v \in \mathcal{V}^{\mathrm{b}}$.
$\ell^{\mathrm{b}}(a) = d_a^{\mathrm{b}} \cdot g(w_a^{\mathrm{b}})$	Effective cost on cycling arc a used in shortest-path computations.
$s_{\mathcal{G}^{\mathrm{b}}}$	Global (average) node strength on $\mathcal{G}^{\mathrm{b}}$.
$L(\mathcal{G}^{\mathrm{b}})$	Line graph of the cycling network.
$E^{\mathrm{b}} = V\big(L(\mathcal{G}^{\mathrm{b}})\big)$	Node set of the line graph (one node per arc of $\mathcal{G}^{\mathrm{b}}$).
$c_{\mathcal{G}^{\mathrm{b}}}$	Global (average) arc-based closeness centrality on $\mathcal{G}^{\mathrm{b}}$.
$l_{\mathcal{G}^{\mathrm{b}}}$	Largest connected component metric.
Ω	Feasible set of upgraded DoS configurations.
B	Available budget for cycling-network upgrades.
$\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$	Cost per metre of upgrading cycling arc a from w_a^{b} to $\tilde{w}_a^{\mathrm{b}}$.
u, v	Origin/destination nodes (in the relevant layer).
$q_{(u,v)}^{\mathrm{b}}$	Cycling demand from u to v (e.g., bikes/hour).
$p_{(u,v)}^{\mathrm{b}}$	AoN shortest cycling path between u and v.
$\mathbb{I}_{a \in p_{(u,v)}^{\mathrm{b}}}$	Indicator: equals 1 if arc a belongs to $p_{(u,v)}^{\mathrm{b}}$, and 0 otherwise.
$\measuredangle_a^{\mathrm{b}}$	Angular change associated with cycling arc a (degrees).
s_a^{b}	Slope factor of cycling arc a.
f_a^{b}	Traffic factor on cycling arc a (derived from DoS and adjacent traffic).
ρ, κ, ψ	Calibration parameters (traffic sensitivity, turn penalty, scaling).
$\mathrm{AADT}(a)$	Annual Average Daily Traffic adjacent to cycling arc a.
$\bar{d}_{(u,v)}^{\mathrm{b}}$	Cyclist-specific generalized distance between u and v.
ϕ_a^{b}	Bicycle flow on cycling arc a induced by AoN assignment.
$\mathcal{E}_a$	Cyclist exposure on cycling arc a (typically $\phi_a^{\mathrm{b}} \cdot f_a^{\mathrm{b}} \cdot d_a^{\mathrm{b}}$).
$\mathcal{E}$	Total cyclist exposure (sum of arc exposures).

$q^{\mathrm{c}}_{(u,v)}$	Car demand from u to v (cars/hour).
ϕ^{c}_a	Car flow on car arc a under user equilibrium.
$t^{\mathrm{c}}_a(\cdot)$	Car travel-time function on arc a (e.g., BPR).
t^0_a	Free-flow travel time on car arc a.
θ^{c}_a	Capacity of car arc a.
$\tilde{\theta}^{\mathrm{c}}_a$	Updated car capacity after reallocating space for cycling upgrades.
μ, σ	BPR parameters (typically $\mu = 0.15$, $\sigma = 4$).
$\mathcal{TTT}^{\mathrm{c}}$	Total car travel time $\sum_{a\in\mathcal{A}^{\mathrm{c}}} \phi^{\mathrm{c}}_a \cdot t^{\mathrm{c}}_a(\phi^{\mathrm{c}}_a)$.
$\mathcal{TTT}^{\mathrm{b}}$	Total bike travel time.
S_{thr}	Safety threshold used in sufficientarian constraints/objectives.
T_{thr}	Time (accessibility) threshold used in sufficientarian constraints/objectives.
$\boldsymbol{w}^{\mathrm{b}}$, $\tilde{\boldsymbol{w}}^{\mathrm{b}}$	Original and upgraded DoS configurations on the cycling network.
$\pi(w^{\mathrm{b}}_a, \tilde{w}^{\mathrm{b}}_a)$	Upgrade cost for cycling arc a (e.g., euros per metre).
$\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}$	Predicted total system travel time (as defined in the chapter).
$\bar{D}^{\mathrm{b}}_a$	Average cyclist-distance measure (as defined in the chapter).
$J_{\mathcal{SI}}$	Total safety insufficiency (relative to S_{thr}).
$J^{\mathrm{b}}_{\mathcal{TI}}$	Total bike-time insufficiency (relative to T_{thr}).
$J^{\mathrm{c}}_{\mathcal{TI}}$	Total car-time insufficiency (relative to T_{thr}).

Chapter 1

Introduction

1.1 Cycling for Healthy and Accessible Urban Mobility

Cycling can play a major role in making cities healthier and more accessible. As a form of active travel, it integrates physical activity into daily routines and is therefore strongly connected to population health. At the same time, cycling provides a low-cost and space-efficient mobility option for short urban trips, improving access to jobs, education, and essential services, especially for people who do not own a car or cannot drive.

Despite these advantages, everyday cycling remains limited in many cities. A central barrier is that many people perceive riding near motor traffic as unsafe or uncomfortable. This makes the quality and continuity of cycling infrastructure crucial: if the network does not offer connected, low-stress routes, potential cyclists often avoid cycling altogether, and the health and accessibility benefits remain unrealized.

In this broader context, cycling can also contribute to environmental and congestion objectives. Bicycles produce no tailpipe emissions, and when they replace motorized trips they can reduce energy use and local air pollution; cycling-oriented policies may also have secondary effects on congestion through mode shift. However, quantifying these system-wide impacts in monetary terms is beyond the scope of this thesis. Instead, we focus on the mechanisms most directly linked to infrastructure design and cyclist experience, namely safety and comfort (through separation from motor traffic and exposure-based indicators) and accessibility (through network continuity and travel-time-based measures) and we study how upgrades shape route choice, flow redistribution, and, in later chapters, mode choice.

The following section synthesizes key evidence on the health, accessibility, and socioeconomic benefits associated with cycling, and motivates the need for planning tools that prioritize investments which improve comfort and reduce stress at the network level.

1.2 Benefits of Cycling

1.2.1 Health Benefits

A large body of evidence synthesizes the health advantages of bicycling and other forms of active travel. Broad reviews conclude that cycling contributes meaningfully to population health across multiple outcomes, from cardiometabolic risk reduction to overall longevity [11, 12]. At the individual level, moderate-intensity activity integrated into daily routines, especially walking or cycling for transport has been linked to healthier body composition in middle-aged adults, suggesting that feasible, everyday activity may be easier to promote and still effective compared with vigorous exercise [13]. Group evidence further indicates that occupational, commuting, or leisure-time physical activity independently lowers the risk of type 2 diabetes in mid-life populations [14]. Daily life energy expenditure shows an inverse, graded association with incident hypertension after adjustment for conventional risk factors, implying benefits across the normotension spectrum and in men at both low and high baseline risk [15]. Domain-specific analyses also report favorable lipid profiles as lower body fat among those who are physically active in commuting or leisure contexts, with additional reductions in diastolic blood pressure observed particularly among men [16]. In line with these findings, commuting by car is associated with a higher risk of first heart attack compared with commuting by bus, bicycle, or walking, whereas higher leisure-time or moderate occupational activity appears cardioprotective after multivariable adjustment [17]. Together, these studies support the importance of distinguishing activity by domain as work, transport, and leisure because each relates differently to specific health indicators and thus merits separate surveillance and policy attention [18].

At the population level, active transportation correlates with a lower prevalence of obesity in all countries. Comparative analyses show that nations with greater distances traveled on foot and by bicycle (e.g., many European countries) have significantly lower obesity rates than the United States, Australia, and Canada, consistent with an inverse relationship between active transportation and obesity on a large scale [19]. Within regions, commuting patterns also matter: in a large Swedish survey, walking or cycling to work and, for men, use of public transport was associated with substantially lower odds of overweight and obesity compared with car commuting [20], and also link active commuting to better fitness and lower body mass index, triglycerides, blood pressure, and insulin (especially in men), reinforcing the potential of daily transport-related activity as a viable health intervention [21].

Health impact assessment studies extend these observational results by estimating net health effects of mode shifts toward walking and cycling. A systematic review of these studies reports that, across diverse settings and modeling approaches, the benefits of increased physical activity generally outweigh the harms from air pollution and traffic injuries, yielding positive benefit–risk or benefit–cost ratios in most scenarios [22]. City-scale analyses arrive at similar conclusions: in Barcelona, substituting a substantial share of long car trips with public transport and cycling was projected to reduce incident cases of diabetes, cardiovascular disease, dementia, and certain cancers in the general population [23]. In urban areas

of England and Wales, a framework scenario for modeling for health improvements has been proposed that suggests a considerable reduction in the burden of disease, with the greatest health benefits attributable to increased physical activity, followed by reduced traffic accidents and air pollution, the results point to solid net benefits from shifting away from private car use for travel [24].

Taken together this evidence, consistently indicates that integrating cycling into daily life improves cardiometabolic profiles, reduces incident chronic disease and premature mortality risks, and can deliver substantial population-level health gains when supported at scale. These findings provide a strong public health rationale for promoting cycling as everyday transport and, in the context of this thesis, for prioritizing network designs and policies that lower stress and increase comfort so that more people can safely realize these benefits.

1.2.2 Accessibility and Equity Benefits

Beyond health, cycling is increasingly framed as an accessibility mode: it can expand the set of opportunities that individuals can reach within reasonable time and comfort, and it can do so with relatively low financial and spatial requirements. From a planning perspective, accessibility is also closely linked to equity, because mobility resources and burdens are not distributed uniformly across populations. Comprehensive transport equity analysis distinguishes between horizontal equity (treating people with comparable needs similarly) and vertical equity (prioritizing disadvantaged groups), and it emphasizes that equity should be assessed both in inputs (e.g., funding, road space, priority) and outcomes (e.g., convenience, comfort, safety, and accessibility) [25]. Since automobile travel is costly, exclusive, and imposes externalities, planning practices that systematically favor driving can widen disparities between drivers and non-drivers. In this sense, a sustainable transport hierarchy that prioritizes resource-efficient modes can support both efficiency and equity objectives [25].

A practical way to connect cycling infrastructure to accessibility is to evaluate whether network improvements reduce stress and expand the reachable set of everyday destinations. For example, [26] proposes a project prioritization method that links reductions in Bicycle Level of Traffic Stress (LTS) to accessibility gains to supermarkets, pharmacies, banks, and libraries, and then examines how these gains distribute across disadvantaged neighborhoods. Their results show that accessibility-based prioritization can highlight projects beyond centrally located corridors, notably along high-stress arterials serving neighborhoods with higher poverty rates, lower vehicle ownership, and legacies of segregation [26]. This illustrates that "where to build" decisions can materially change who benefits from cycling investments.

Accessibility also extends naturally to multi-modal contexts. Bike–transit integration can increase access for both transit users and cyclists by enabling longer, more connected trips. Using high-resolution transit data, [27] analyzes bike–transit accessibility increments and shows that integration can substantially expand reachable opportunities, while also revealing inequities: some communities experience lower accessibility increments and lower multimodal ridership despite high transit dependence, suggesting that targeted cycling infrastructure im-

provements can address accessibility gaps [27]. At a more behavioral level, perceived accessibility is influenced by individual characteristics and neighborhood context. Evidence from dockless bike-sharing indicates that cycling can support participation in daily activities (particularly commuting and access to transit), and that the availability of cycling paths in the home neighborhood is positively associated with perceived accessibility benefits [28].

Taken together, this literature motivates evaluating cycling investments not only through aggregate network performance, but also through accessibility outcomes and their distribution across groups and neighborhoods. In the context of this thesis, this supports (i) focusing on network continuity and low-stress routing as prerequisites for meaningful accessibility gains, and (ii) incorporating distribution-sensitive objectives in later chapters to assess whether upgrades help ensure minimum standards of safety and access for a broad range of users.

1.2.3 Economic and Environmental Benefits

Cycling delivers value through multiple, reinforcing channels. At the system level, it lowers external costs from congestion, greenhouse-gas emissions, air pollution, and noise; at the individual level, it improves health and productivity; and at the regional level, it stimulates employment across repair, tourism, services, and infrastructure. These channels are now quantified in transport and cost-benefit assessment frameworks.

A growing body of work monetizes health benefits using the WHO's Health Economic Assessment Tool (HEAT), which translates reduced all-cause mortality from active travel into economic terms [29]. Applied to a proposed canal-side route near Dublin, HEAT indicates substantial net benefits and strong benefit–cost ratios for realistic mode-shift scenarios [30]. At national scale, Dutch cycling levels (roughly three-quarters of an hour per adult per week on average) are associated with about 6,500 prevented deaths per year and €19 billion/year in health-economic value [31]. These magnitudes illustrate why even moderate increases in cycling can have outsized fiscal and societal returns.

Beyond health, cycling compares very favorably to driving when full social costs are considered. In 2015 Copenhagen's cost–benefit accounting attributes about €0.08/km in social cost to cycling versus €0.50/km to car use, once travel time, safety, emissions, noise, and health are all included [32]. Although time is the main private cost for cyclists (lower operating speed), the trend in total social cost favors bicycle use as urban congestion and climate externalities increase.

Macro-scale scenarios reinforce these ideas. For the United Kingdom, two political objectives doubling the level of bicycle travel by 2025 (Cycling Delivery Plan, CDP) and reaching 10% by 2025 and 25% by 2050 (Get Britain Cycling, GBC)—were evaluated using a simulated population and standard monetization of health, amenity, decongestion, safety, noise, and air-quality effects [33]. The GBC's greater ambition generates much higher annual benefits by 2050, with decongestion dominating external savings and additional gains from reduced emissions, quieter streets, and cleaner air. Quantitatively, the CDP pathway is associated with about £2.09 billion per year by 2025 (roughly £36 per person) rising to £6.41 billion by 2050

(about £98 per person). Under the more ambitious GBC pathway, annual benefits are already around £6.37 billion by 2025 (about £109 per person) and reach approximately £42.63 billion by 2050 (about £649 per person).

Breaking down the external components clarifies where the gains accrue [33]. Decongestion is the largest single contributor in every case: on the CDP pathway, it amounts to roughly £0.28 billion per year by 2025 and £0.96 billion by 2050; for GBC, the decongestion effect is about £1.09 billion in 2025 and £7.09 billion by 2050. Emissions (greenhouse gases) reductions add materially—on the order of £10.9 million (CDP 2025) to £36.6 million (CDP 2050), and from £41.6 million (GBC 2025) to £271.6 million (GBC 2050). Noise abatement contributes several million per year (roughly £2.4–£8.0 million under CDP and £9.1–£59.2 million under GBC, for 2025 and 2050 respectively), while air-quality improvements add a similar magnitude (about £1.5–£5.2 million for CDP and £5.9–£38.6 million for GBC across 2025–2050). So the pattern that emerges is that decongestion leads by a wide margin, with climatic, atmospheric, and acoustic benefits adding to the total.

Cycling also contributes to local economies through employment and inclusive growth. Job creation per unit of turnover is typically higher in cycling than in many other transport activities, particularly with regard to infrastructure construction, while also providing accessible jobs for less-skilled workers in the retail, repair, tourism, and service sectors [34]. At the workplace level, cycling to work is consistently associated with greater well-being and fewer sick days, resulting in a measurable increase in productivity [35, 36].

Overall, the evidence is consistent across all methods (HEAT, CBA, sectoral employment studies, and national scenarios): shifting trips to bicycle use yields large net economic and environmental benefits, with positive effects on health and traffic congestion, as well as improvements in climate, air quality, noise, and productivity, thereby increasing the return on investment.

1.3 Barriers to Cycling

Despite clear health and economic gains, everyday cycling remains constrained by a mix of perceived and objective barriers related to safety, infrastructure continuity, social norms, and trip context [2]. The literature points to three interlocking layers that shape whether people feel both able and willing to ride: the built environment and facility design; social and organizational supports; and individual attitudes and constraints.

A first set of barriers arises from perceived safety and the physical characteristics of routes. Evidence from Portland shows that simple proximity to trails or painted bike lanes is not, by itself, a reliable predictor of cycling; instead, what matters is whether people believe bike lanes are available and whether the street network offers connected, direct paths for utilitarian trips [37]. When cyclists use their bicycles, GPS data reveals that they clearly prefer paths away from public roads, bike lanes in neighborhoods with calm traffic, and well-designed intersections, while conventional bike lanes mainly compensate for the stress of adjacent traffic

rather than adding positive appeal beyond a quiet street with low traffic volume [3]. Distance also influences feasibility: in Montreal, the average distances for walking and cycling trips were approximately 650 m and 2 km, with substantial variations depending on the purpose of the trip, urban form, and sociodemographic factors [38]. Furthermore, specific exposures such as the "door zone" on streets with on-street parking maintain a sense of risk; video observations in mixed-use corridors recorded frequent interactions with parked vehicles and occasional door openings in the path of cyclists, underscoring the value of designs that minimize door conflicts when complete separation is not yet possible [4].

The second layer relates to social norms and the support available at destinations, which help turn intention into actual behavior. Workplace culture and facilities such as secure bike parking, storage rules that allow bicycles, showers, and simply seeing colleagues arrive by bike are consistently associated with a higher probability of cycling at least once a week. These effects appear particularly strong for women: even one supportive cultural signal combined with one physical facility can noticeably increase the likelihood of commuting by bike [39]. In university settings, surveys in the UK show that students recognize the enjoyment and health benefits of cycling, but bad weather and safety concerns are the main reasons they do not cycle more. Many also report that cycling would be easier "somewhere else," which underlines the importance of local context and infrastructure quality in closing the gap between positive attitudes and actual behavior [40]. Similar results are found in Spain: active commuting is positively related to perceived self-efficacy and the presence of walking and cycling infrastructure, and negatively related to planning barriers. Overall, the evidence suggests that changes in commuting behaviour mostly reflect a shift between travel modes [41].

A third layer involves attitudes, perceived behavioral control, and local context. In a Brazilian medium-sized city, respondents expressed positive views of cycling's environmental, health, and economic benefits but negative views about effort, speed, prestige, weather/appearance, and safety. The strongest deterrent was inadequate infrastructure, followed by safety concerns and topography. Crucially, respondents with more positive attitudes perceived weaker barriers, suggesting that better facilities and supportive social cues do not merely remove constraints but also shift beliefs about what is practical [42]. In the Central and Eastern European context, a Warsaw study found many residents framing cycling as leisure rather than transport; those who did use the bike for everyday trips emphasized speed and flexibility. Main obstacles were poor facility quality and continuity, driver behavior, and winter maintenance signaling the need for both network design improvements and culture change in traffic norms [43].

These micro-level findings sit alongside macro trends in who currently rides and which policies move the needle. Using national US travel surveys, walking increased slightly from 2001 to 2009 while cycling remained rare, with active travel concentrated among men, middle-aged, employed, higher-educated, and car-free groups. The authors conclude that infrastructure must be coupled with targeted programs for underrepresented users such as women, children, and seniors [44]. Where cities have built protected facilities at scale, safety has improved and participation has broadened. International and city-level analyses document large declines in cyclist fatalities per capita and per distance cycled since the 1990s, with some of the most rapid

gains in places that implemented coherent networks of protected lanes and paths (e.g., Bogotá, Sevilla) [45]. Complementary traffic calming of residential streets (30 km/h limits enforced by geometric design) further expands the low-stress network without requiring a bikeway on every block; such measures are widespread in the Netherlands, Germany, Switzerland, and Austria, but patchier elsewhere [45].

Taken together, the evidence suggests that everyday cycling is constrained by a combination of (i) the physical characteristics and continuity of the built environment, (ii) social and organizational supports at origins and destinations, and (iii) individual attitudes, capabilities, and trip contexts. This thesis does not attempt to address all three layers simultaneously. Instead, we focus on the infrastructure and network-design dimension, because it can be represented explicitly within a graph-based modeling framework and translated into implementable upgrade decisions (e.g., changes in the degree of separation, continuity of low-stress routes, cyclist flow, and exposure to motor traffic). The other dimensions such as social supports and behavioral factors remain important, but are outside the scope of this work; we view them as natural directions for future research, including extensions that incorporate destination facilities, policy incentives, and heterogeneous cyclist preferences alongside infrastructure investments.

1.4 Literature Review

Our work draws on four primary domains: bicycle facility categorizations, cycling network optimization, and analyses of modal shift and transport justice. This subsection reviews the main street and facility classification frameworks that connect roadway attributes to cyclist stress, comfort, and safety.

1.4.1 Bike Lane and Street Classification Frameworks

A natural starting point is the user perspective. Geller's typology [8] divides the adult population into four groups by tolerance for traffic stress: Strong and fearless, Enthused and confident, Interested but concerned, and no way, no how. In this context, traffic stress is a behavioral notion: it captures whether cycling on a link feels acceptable given exposure to motor vehicles (speed/volume), the degree of separation, and conflict points such as intersections. The LTS approach formalizes this by assigning each segment (and intersection crossing) to a small number of stress levels, allowing planners to measure "low-stress connectivity" for typical users rather than only for confident riders. This motivates classifications that encode stress/comfort explicitly and supports network design for the interested-but-concerned majority.

Early quantitative approaches focused on roadway segments and combined traffic exposure with geometry. Davis's Bicycle Safety Index Rating (BSIR) [46] introduced a two-part structure: a Roadway Segment Index for midblock conditions and an Intersection Evalua-

tion Index for nodes. Core inputs include average daily traffic, number of lanes, posted speed, width of the outside through lane, pavement condition, and a location or context factor reflecting on-street parking, curves, grades, and driveway activity. The combined score is mapped to qualitative categories from excellent to poor, giving practitioners a way to compare corridors with differing volumes, speeds, and widths while penalizing poor surface quality and complicated contexts.

Building on BSIR, planners in Hollywood, Florida, proposed the Epperson–Davis Roadway Condition Index (RCI) [47]. RCI simplified the assessment by dropping the intersection component to concentrate on midblock riding conditions and by reducing the influence of pavement and location factors that had been relatively strong in BSIR. A key change is the interaction of geometry and operating speed: lane width is multiplied by speed limit to capture the much greater stress of narrow lanes at higher speeds. Inputs continue to include average daily traffic, lane count, outside lane width, speed, and basic surface/context checks, but weight is shifted to reflect the risk of narrow, fast-running cross-sections.

Landis's Interaction Hazard Score (IHS) [48] extends segment-based scoring with empirical calibration. The model integrates average daily traffic, total through lanes, usable width of the outside through lane including any marked bike lane, speed limit, the proportion of heavy vehicles, a pavement factor derived from a five-point surface rating, adjacent land-use intensity, and the frequency of uncontrolled access points such as curb cuts and parking turnover. Unlike purely formulaic approaches, IHS coefficients were tuned against rider surveys, so variables such as trucks, driveway density, and dooring exposure receive influence that better matches perceived comfort rather than relying only on volume and speed.

In parallel, Sorton and Walsh proposed the Bicycle Stress Level (BSL) [47], which offers a transparent, additive measure with three primary drivers: motor traffic volume, effective lane width, and vehicle speed. The score can be adjusted for commercial driveway density, parking turnover, and heavy vehicle share. Results are grouped from 1 (very low stress) to 5 (very high stress). The clarity of the three main levers makes BSL useful for communicating how widening a lane, lowering speeds, or reducing exposure can move a segment into a more comfortable category.

Turner's Bicycle Suitability Score (BSS) [49] organizes Texas roadways into categories by checking geometric and traffic conditions and whether any bicycle facility is present. Inputs include traffic volume, prevailing speed, pavement width, and the presence of bike lanes, paved shoulders, or parallel off-street paths and sidewalks. The output is a five-class typology that distinguishes street-oriented routes with bicycle facilities, low-volume alternatives without formal markings, links without in-street facilities but with sidewalks, links with neither bicycle nor pedestrian facilities, and off-street shared-use paths. Residential low-speed streets are acknowledged as broadly suitable even though frequent intersections limit uninterrupted travel.

The Federal Highway Administration's Bicycle Compatibility Index (BCI) [50] consolidates many of the above ingredients and anchors them in perception data. Inputs include the presence and width of a bike lane or shoulder, curb-lane width, traffic volumes in curb and adjacent lanes, 85th-percentile speed, the presence of active parking, roadside development

type (residential or other), and an adjustment factor accounting for trucks, parking turnover, and right-turn movements. Lower BCI values correspond to greater comfort. To aid design decisions, BCI maps the numeric score to a six-level service scale, so engineers can relate a specific cross-section to an expected comfort band for typical adult riders.

Bicycle Level of Service (BLOS), incorporated into the 2010 Highway Capacity Manual [51], standardizes an A–F rating for shared roadways using a set of logarithmic forms. The model uses short-interval directional motor traffic volume, total through lanes, posted speed as a proxy for actual running speed, percentage of heavy vehicles, commercial trip-generation intensity of adjacent land uses, the effective frequency of uncontrolled accesses such as driveways per mile, a five-point pavement condition rating, and the effective outside-lane width combining pavement, shoulders, and encroachments. BLOS advanced consistency and comparability across studies and jurisdictions, though the functional form and calibration constants can make individual contributions less intuitive to non-specialists.

A shift toward user inclusivity and network continuity is represented by the Level of Traffic Stress (LTS) framework by Furth, Mekuria, and Nixon [52]. LTS assigns each link and intersection crossing a stress level from 1 to 4 based on simple, defensible thresholds for prevailing speed environment, lane count, presence and width of bike lanes or protected tracks, presence and configuration of parking and associated door zone risk, and median or buffer separation. LTS 1 and 2 identify low-stress links suitable for the interested-but-concerned majority and align with Dutch practice, LTS 3 reflects conditions comfortable mainly for experienced riders, and LTS 4 captures high-stress segments that most people will avoid. By extending the logic from links to intersections, LTS supports network-level gap analysis and route building that prioritizes continuity of low-stress travel.

Building on these approaches, Campero-Jurado et al. [53] propose the *Degree of Separation* (DoS), a tag-driven categorization inspired by Geovelo[1] that groups cycling facilities into six ordered classes according to the strength of physical or operational separation from motor traffic [54]. The ordering follows a consistent empirical finding: separation is one of the main drivers of perceived comfort and safety [55]. In practice, DoS classes are inferred from OpenStreetMap tags capturing segregation (e.g., cycleway=track), access control (e.g., motor-vehicle restrictions), speed environment and configuration (e.g., maxspeed, speedlimit), as well as context factors such as pedestrian sharing or parking, with surface-related tags used to exclude poor-quality segments. The resulting hierarchy is designed to be monotone in protection and, in our data, is consistent with independent safety/comfort scoring at the facility-type level, see Table 1.1.

DoS complements earlier segment-based assessment methods (e.g., BSIR, BCI/BLOS) and thresholded stress frameworks such as LTS by offering a separation-first classification that is simple, scalable, and directly extractable from OpenStreetMap. We adopt DoS as the backbone representation in this thesis because it is OSM-native (and therefore reproducible and transferable across cities), easy to communicate, and naturally supports upgrade modeling as monotone improvements in link quality.

[1] `https://gitlab.com/geovelo-public/requetes_amenagements_cyclables/-/tree/master`

A key limitation, however, is that OSM tags alone do not reliably capture several determinants highlighted in the literature, particularly traffic conditions and detailed cross-section features (e.g., parking interaction, cars flow), which may be missing or inconsistent in open data. To retain DoS as the decision variable while still reflecting these factors in the network weights, we align DoS levels with the CycleRAP Safety and Comfort Score (CSCS) [5]. CycleRAP is an internationally applicable scoring framework that rates cycling links based on observable infrastructure and traffic related characteristics such as separation from motor traffic, speed environment, traffic volume and composition (including heavy vehicles), and curb-side risks, thereby providing a consistent way to approximate safety and comfort beyond facility type alone. Through this alignment, our arc weights inherit CSCS informed penalties, so that optimization and routing reflect not only the DoS category but also the broader set of design and traffic contributors captured by CycleRAP.

Table 1.2 summarizes how different classification approaches account for these factors. In particular, it highlights that while DoS is primarily an OSM-derived separation measure, combining it with CSCS effectively expands the set of characteristics represented in our model. This is why we use DoS as the network classification and rely on CSCS to calibrate the corresponding link weights within a unified and transferable framework. To clarify the ordering of intermediate classes (DoS 2–4), note that these levels do not differ only by whether a bicycle space is marked, but also by the traffic environment typically associated with the corresponding OSM tags. DoS 4 mainly captures low-speed, low-volume streets (e.g., residential/traffic-calmed links) where cycling is allowed and mixing with cars can remain relatively comfortable because interactions with motor traffic are limited. DoS 3 represents shared on-street situations such as advisory markings or pictograms, which often occur on busier streets and therefore involve more frequent interactions with motor vehicles despite the presence of markings. DoS 2 corresponds to painted bike lanes adjacent to general traffic; although they allocate a dedicated lateral space, they are frequently located on higher-speed or higher-flow corridors and may coincide with curbside parking, where close passing and dooring risk can dominate perceived comfort. The CSCS calibration accounts for these contextual effects (traffic intensity, speed environment, and adjacent parking exposure), which explains why a painted lane (DoS 2) can receive a lower safety/comfort weight than a calm mixed-traffic street (DoS 4).

DoS	Key OSM tags (indicative)	Typical facility	CSCS	AADT	Mean speed	Interpretation
6	`highway=cycleway`; `cycleway=track`; `segregated=yes`	Protected cycle tracks / dedicated cycleways	5.23	14 400	50 km/h	Strong separation despite higher traffic
5	`highway=path|footway` + `bicycle=designated`; `motor_vehicle=no`	Greenways / shared or segregated paths	4.64	16 000	50 km/h	Separated from cars, possible pedestrian interactions
4	`maxspeed=30`; `living_street|pedestrian` (cycling allowed)	Traffic-calmed mixed streets (30 km/h)	2.99	960	30 km/h	Low speed, but mixing with cars
3	`cycleway=shared_lane`; `oneway:bicycle=no`	Sharrows / advisory / contraflow without protection	2.18	3 120	30 km/h	Markings only; comfort depends on traffic
2	`cycleway=lane|shoulder`	Painted bike lanes on traffic streets	1.55	6 912	40	Paint-only separation; speed/volume matter
1	`cyclestreet=yes` or `bicycle_road=yes`; `cycleway=share_carway`	Unmarked/shared lanes; context-dependent links	1.00	8 448	40 km/h	No explicit protection; bike and cars share road

Table 1.1: DoS levels derived from OSM tags and their alignment with the CycleRAP Safety and Comfort Score (CSCS), illustrated with Grenoble traffic context (AADT and mean speed near cycling links). Higher DoS indicates stronger separation from motor traffic and corresponds to higher CSCS values.

Methodology	Traffic Volume per Lane	Curb/Lane Width	Vehicle Speed	Speed Limit	Pavement Type/Condition	Parking/Turn Lanes	Grades	Sight Distance/Visibility	Driveway Frequency	Adjacent Land Use	Signalization/Intersections	Heavy Vehicles	Vehicle LOS	Maintenance	TDM / Multimodal Support	Provision / Type of Facility	Cyclist Input	Passing / Meeting Freq.	Service Volumes	Bicycle Density	Curve Radius	Total Delay	Avg. Bicycle Speed
(col.)	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**
BSIR (Davis, 1987)	✓	✓		✓	✓	✓	✓		✓		✓										✓		
Epperson–Davis RCI	✓	✓		✓	✓	✓	✓	✓	✓												✓		
IHS (Landis, 1994)	✓	✓		✓	✓	✓			✓	✓		✓					✓						
BSL (Sorton & Walsh, 1994)	✓	✓	✓			✓			✓			✓											
BSS (Turner, 1997)	✓	✓	✓													✓							
BCI (Harkey et al., 1998)	✓	✓	✓			✓				✓		✓				✓							
BLOS (HCM 2010 / Landis)	✓	✓		✓	✓	✓			✓	✓		✓											
LTS (Furth et al., 2016)		✓	✓	✓		✓					✓					✓							
DoS (Campero Jurado et al., 2024)	✓	✓	✓	✓	✓		✓				✓	✓		✓		✓					✓		✓

Table 1.2: Unified summary of input variables used by bicycle facility assessment methods. A checkmark indicates explicit use of the variable in the method (columns 1–12 correspond to roadway/traffic and context variables; columns 13–23 cover additional operational and service variables).

1.4.2 Bike Network Optimization and Expansion

A first line of work treats the design of cycling infrastructure as an explicit optimization problem in a given road network. In [56], the introduction of on-street bike lanes is formulated as a bilevel problem: the upper level chooses which car lanes to reallocate to bikes, trading off total travel time for motorists against the amount of cycling that can be routed on dedicated space, while the lower level assigns car and bike flows. The bilevel model is solved with a genetic algorithm, and numerical experiments illustrate how different weightings between car and bike objectives lead to different network designs. A related approach in [57] optimizes where to add bike routes and lanes in an existing urban network, focusing on connecting origin–destination pairs with routes that are short and have a good bicycle level of service, illustrated with a case study in Baltimore. The bikeway network design model in [58] goes further and formulates a multi-objective grey 0–1 program that minimizes cyclist risk and impacts on motor traffic while maximizing comfort and service coverage under budget and continuity constraints. A district-scale case study in Taipei shows how uncertain inputs can be handled through grey numbers and how the model generates several non-dominated network alternatives rather than a single plan. In [59], a mixed-integer multi-commodity flow model is proposed in which cyclists follow shortest paths on a partially built network; links without cycling infrastructure incur higher user costs and discontinuities are explicitly represented. Because exact solution becomes expensive on real street graphs, a metaheuristic is developed that scales to Montevideo and can additionally minimize the number of network breaks.

A second group of contributions optimizes not only where to build, but also how and when to expand networks or related systems. In [60], continuous approximation is used to derive optimal spacing of bike lanes, bike-sharing station locations and fleet size in cities with heterogeneous demand, balancing user costs (access and in-vehicle time) against agency costs. A case study in Guayaquil in Ecuador, shows that user costs dominate total costs and that higher demand can often be accommodated by densifying the network rather than scaling facilities proportionally. The long-term expansion of cycle superhighways over a 50-year horizon is studied in [61]. There, monetized benefits from a regional transport model are mapped back onto individual links and integer programs decide which segments to build when, approximating travel time benefits linearly but capturing cumulative effects. The optimized rollout schedule yields a much higher net present value and benefit–cost ratio than the actually implemented investment sequence. In [62], the focus shifts to the joint allocation of road space between bikes and cars, through a linear programming formulation that searches for Pareto-efficient trade-offs. For a given city and budget, the model produces families of scenarios that reassign lanes to cycling while explicitly tracking the consequences for car accessibility, providing planners with a menu of reallocation options rather than a single “best” solution.

Data-driven approaches form a third major cluster, using large-scale empirical information on trips, trajectories or crashes to guide network growth. In [63], bicycle networks in many cities are shown to consist of numerous disconnected components; greedy algorithms are then introduced that add links to maximize connectivity and directness of the bicycle layer using OpenStreetMap data. With only a small number of carefully chosen new links, connectedness improves substantially, and the algorithms outperform random or purely local strategies. The

framework in [64] couples mobile-phone based origin–destination matrices with GPS traces from a bike app to estimate potential cycling demand in Bogotá; percolation theory is then used to prioritize street segments whose conversion to bike paths would both capture high potential trip volumes and preserve global connectivity. In [65], demand is incorporated in reverse: starting from a fully built bike path network, links are iteratively removed while cyclists' route choices are recomputed under a simple safety-based routing model. This adaptive percolation process yields families of networks that remain well adapted to spatially heterogeneous demand even under strong infrastructure reductions, and allows direct comparison between demand-driven and demand-agnostic designs. The Manchester case study in [66] combines distance decay, route calculation and network analysis with census origin–destination data and local topography to prioritize corridors where high cycling potential coincides with the possibility of low-stress routing. The framework explicitly embeds egalitarian criteria, showing that investments can be distributed more fairly across neighborhoods without sacrificing aggregate benefit. In [67], trip and crash data for bikes and e-scooters in Turin are integrated into a planning tool for micromobility. A tunable parameter shifts the optimization between purely demand-driven and purely safety-driven network extensions, and systematic experiments show that intermediate settings can produce better long-term network quality than either extreme. In [54], graph-theoretic metrics of French cycling networks are related to survey-based comfort and safety scores to learn how topology affects perceived conditions; this relationship is then used to define a topology optimization problem that selects new links to maximize predicted comfort under budget constraints, solved with a customized algorithm that bridges network science and planning practice.

Several contributions emphasize that realistic network design must account for route choice and user preferences. In [68], detailed bike-sharing trajectory data are used to build an integer optimization model that chooses where to build lanes so as to maximize cyclists' utility, with route choice represented by a multinomial logit model in a bilevel formulation. The resulting mixed-integer programs can be solved efficiently, and a real-world case study with a city government reveals trade-offs between coverage of trips and continuity of bike lanes. A complementary perspective is proposed in [69], which treats the route choice model itself as unknown and seeks to infer it from observed cyclist movements. The authors assume that cyclists follow shortest paths under a personalized convex combination of several link attributes (e.g., distance, safety, pollution/noise), and that the population is composed of a small number of user groups with different weight vectors. Using link flow observations (potentially on only a subset of arcs), they formulate an inverse optimization problem that jointly identifies representative preference profiles and the share of cyclists associated with each profile. Such identification approaches can provide data-driven inputs for network design, enabling planners to tailor upgrades to heterogeneous cyclist preferences rather than relying on a single "average" bicycle level-of-service formulation.

A final line of work focuses more on strategic narratives and policy framing than on specific optimization models. The think piece in [70] argues that current transport policy is trapped between contradictory goals and incrementalism, and proposes the idea of an "E-Bike City" as a starting point for rethinking urban mobility around cheap, flexible active modes rather than private cars. This agenda calls for integrated research that connects infrastructure

optimization, emerging vehicle types and broader regulatory tools.

Across these contributions, there is a clear progression from early deterministic network design models towards rich, data-driven and multi-criteria formulations. Some approaches treat bike lanes as a stand-alone layer to be optimized given exogenous demand and car traffic; others co-optimize bikes and cars, or let infrastructure co-evolve with demand and route choice. A growing subset uses large-scale digital traces to infer where, how and for whom to build, and some explicitly incorporate social equity and long-term rollout timing. Our work positions itself in this space by combining network-scale optimization with comfort- and safety-oriented metrics that can be derived from open data, aiming to provide planners with tools that are both analytically grounded and easily transferable to cities where only basic map and demand information are available.

Table 1.3 summarizes the main methods reviewed above, indicating for each study the key data inputs and the primary optimization objectives.

Paper	Method	Data	Objective (sketch)
Paulsen & Rich (2023) [61]	Integer programming; time-staged expansion	Regional bike network; benefit mapping from CBA	Max net present value; optimal timing/prioritization of links
Mesbah et al. (2012) [56]	Bilevel + genetic algorithm	Synthetic/road network; traffic and bike demand	Trade-off: minimize car time vs. maximize bike use on lanes
Duthie & Unnikrishnan (2014) [60]	LP/continuous approximation	Demand patterns; bike-share layout	Minimize total cost; optimize lane spacing/length, stations, fleet
Lin & Yu (2013) [58]	Multi-objective IP	Traffic accident and network data	Maximize safety/comfort and coverage; minimize impacts on other modes
Mauttone et al. (2017) [59]	Multi-commodity flow MIP	OD survey data	Cost–benefit trade-off with route coverage constraints
Liu et al. (2022) [68]	MILP	Bike-share trajectories	Maximize coverage of likely route choices under budget
Steinacker et al. (2022) [65]	Heuristic (betweenness)	Bike-share + OSM	Minimize bike travel times; prioritize central links
Ballo & Axhausen (2024) [70]	Heuristic (betweenness)	Census + OSM	Reduce impact on car network while enhancing bike network
Natera Orozco et al. (2020) [63]	Heuristic (greedy)	OSM only	Improve network connectedness with incremental link additions
Olmos et al. (2020) [64]	Percolation theory	Mobile-phone + census	Maximize demand coverage via connectivity growth
Mahfouz et al. (2023) [66]	Route calc. + network analysis	Distance decay; street network; equity targets	Prioritize low-stress yet direct links under budget/equity scenarios
Folco et al. (2023) [67]	Data-driven prioritization	Micromobility trips + crashes	Tune demand vs. safety emphases; improve long-run network quality
Jurado et al. (2024) [54]	Graph-metric prediction + budgeted optimization	OSM network + surveys	Maximize predicted perceived comfort/safety under budget
Ardizzoni et al. (2024) [69]	Optimization-based inverse shortest-path estimation; convex programming + clustering of user groups	Observed cyclist trajectories / link-level flows	Infer multiple user "types" by identifying their weight vectors over basic road attributes, estimate the probability of each user type, and recover the implicit multi-criteria BLOS perceived by different cyclist groups.
Wiedemann et al. (2025) [62]	Linear programming; Pareto frontier	Real/synthetic networks	Lane reallocation scenarios; balance bikeability and car accessibility

Table 1.3: Representative optimization/prioritization approaches for bike network design and expansion.

1.4.2.1 Bike Mode Choice

Utility cycling has become an increasingly credible alternative to driving in urban areas, yet automobility still dominates both street space and policy agendas [71, 72]. Work on the politics of infrastructure shows that decisions about where and how to build cycling facilities are not merely technical but deeply tied to questions of social and spatial justice, marginalization, and competing claims on urban space. Cycling infrastructure is expected to protect users from the very real dangers of motor traffic, to signal that cyclists are legitimate road users, and at the same time to challenge car dominance rather than simply accommodating it within the existing paradigm [71].

A large empirical literature links these political aspirations to concrete design choices. Reviews of bikeway research show a consistent hierarchy in user preferences: people favour separated paths and protected lanes over riding in mixed traffic, and poor pavement quality is a strong deterrent [73]. Intersection design is highlighted as a recurrent weak point, since junctions amplify perceived risk even when mid-block facilities are good. Overall, the evidence suggests that coherent networks of low-stress links and crossings are more important for encouraging cycling than isolated facilities, and that these networks need to serve diverse users and trip purposes rather than only confident riders [72, 73].

Behavioral and modeling studies complement this infrastructure perspective by examining how people actually choose modes and routes. An agent-based modal choice model in [74] represents individuals who evaluate alternatives along multiple criteria (time, comfort, cost) but whose decisions are shaped by habits and cognitive biases. Calibrated with survey data and implemented in a simulator, this framework shows that factors such as status quo bias or misperception of travel times can blunt the impact of new cycling infrastructure or financial incentives, and that policy scenarios which look effective in a purely rational model may have much smaller effects once habits are taken into account. Related computational work uses smartphone traces and app-based data to estimate door-to-door travel times and emissions for different modes, and to quantify how many existing trips could shift to cycling or other low-carbon options under realistic conditions [75, 76].

The potential for such a shift has been examined more explicitly using urban travel surveys. A methodology developed in [77, 78] identifies which car trips in a city can realistically be transferred to walking, cycling, or public transport without worsening travel time and while respecting trip purpose and other constraints. In an application to Madrid, only a minority of car trips met these conditions, but the method showed clear spatial patterns: central areas had higher potential for transfer to walking and cycling, while peripheral zones favoured public transport. The studies underline that even in dense European cities, the substitution potential is constrained by trip length, activity patterns and current network performance, and that infrastructure improvements must be combined with broader policy measures if substantial modal shift away from the car is to be achieved.

Taken together, this body of work shows that bike mode choice emerges from the interaction of three elements: the political framing of cycling within transport policy, the physical

characteristics and continuity of cycling networks, and the behavioural processes by which individuals evaluate and adapt their travel habits. Most existing models either treat infrastructure quality in coarse terms or do not distinguish clearly between different safety or comfort levels of cycling facilities. This motivates approaches that connect fine-grained, safety-oriented infrastructure classifications to explicit mode-choice mechanisms, in order to better estimate the impact of specific upgrades on the shift from cars to bicycles.

1.4.3 Transport Justice

Traditional transport appraisal has largely followed a utilitarian logic [79, 80], seeking to maximize aggregate welfare (for instance by minimizing total travel time) while treating the distribution of benefits and burdens as secondary. Inequalities are acceptable as long as the total sum improves. In political philosophy, Rawls offers a contrasting view: justice requires both equal basic liberties and a fair distribution of social and economic advantages, where inequalities are justified only if they benefit the least advantaged and positions remain genuinely open to all [81].

Building on this, [7] recasts transport planning around accessibility rather than mobility, and around justice rather than efficiency. Accessibility is treated as a primary social good, and the central question becomes not "how can we optimize flows?" but "do all individuals reach at least a reasonable minimum of opportunities?" This sufficientarian perspective implies explicit accessibility thresholds that everyone should attain and requires planners to prioritize improvements for groups and areas falling below these minima. In doing so, it offers a concrete operational reading of Rawlsian and related egalitarian ideas in the transport domain and directly challenges conventional practice that tolerates large accessibility gaps as long as average performance improves.

Recent work begins to translate these justice principles into optimization models and quantitative tools, [9] formulate the operation of an intermodal autonomous mobility-on-demand system with the explicit objective of minimizing accessibility unfairness rather than total travel time. Using a network-flow representation and a fairness criterion defined over origin–destination pairs, they show that a fully fair solution (in terms of average accessibility) can be reached at the cost of only a modest increase in aggregate travel time, thereby demonstrating that putting justice at the center of system design is not only normatively desirable but also operationally feasible. In parallel, [82] compare different justice theories—utilitarianism, egalitarianism, communitarianism, and the capabilities approach—in the equity assessment of a metro expansion in Lisbon. By constructing accessibility indicators consistent with each theory, they show how the choice of normative lens changes which groups are seen as winners or losers and argue for minimum accessibility standards informed by a capabilities perspective.

Overall, this literature marks a shift from viewing equity as an ex-post diagnostic to treating justice as a primary design criterion. Martens' accessibility thresholds and the optimization-based framework of [9] are particularly relevant for our work: they show how transport systems can be planned and operated around explicit fairness targets, a perspective

we extend to the safety and comfort of cycling networks.

1.5 Thesis Motivation and Contributions

1.5.1 Motivation

Cycling can deliver major public benefits, including improved health, reduced emissions, and more livable cities, yet it remains a minority mode in many urban areas because many potential riders perceive cycling in mixed traffic as unsafe or uncomfortable. A consistent finding in the literature is that physical separation from motorized traffic is one of the strongest determinants of perceived safety and comfort, and therefore of cycling uptake. Importantly, the benefits of cycling infrastructure are often non-additive: well-chosen upgrades can create network synergies, where connecting a few critical gaps unlocks continuous low-stress corridors and reroutes many trips, so that the overall benefit is greater than the sum of isolated local improvements (informally, "$1 + 1 > 2$"). These synergies arise precisely because safety, comfort, and usability depend on network continuity, route choice, and (in multi-modal settings) demand and congestion feedback.

However, upgrading infrastructure is costly and competes with other urban priorities. Cities therefore require decision-support tools that can identify where limited budgets should be invested to deliver the largest safety and comfort gains, while avoiding unacceptable negative impacts elsewhere in the transport system. Capturing the above synergies is challenging, and several methodological gaps motivate this thesis. First, many existing approaches assess safety at the level of individual street segments, whereas effective investment decisions depend on network effects such as continuity and the existence of connected low-stress corridors; this calls for network-scale metrics that remain interpretable and actionable. Second, when planning is modeled as upgrading links (increasing their degree of separation), the metrics used to guide optimization should behave monotonic: improving one or more links should not worsen the network-level score. Yet many standard graph metrics do not satisfy such monotonic behavior, which can lead to misleading optimization outcomes. Third, topology alone can misrepresent where upgrades matter most in practice: structurally central links may carry little cycling demand, so their upgrade may benefit few riders, whereas upgrades that shape route choice can yield large, network-wide exposure reductions. Fourth, cycling upgrades often require reallocating road space and thus interact with car traffic performance; ignoring multi-modal feedback can mischaracterize the real trade-offs and the conditions under which synergies appear. Fifth, demand should not be treated as fixed: improving cycling conditions can induce modal shift from car to bicycle, changing both cycling and car flows and potentially mitigating congestion impacts. Sixth, infrastructure planning is inherently distributional: improving average outcomes does not guarantee that minimum standards of safety and accessibility are met for all users or neighborhoods, motivating the inclusion of explicit transport-justice objectives. Finally, cross-city benchmarking remains challenging because networks differ greatly in size, spatial coverage, and mapping resolution; direct comparisons of raw graphs can therefore conflate safety differences with representation effects, calling for

size-robust comparison tools.

1.5.2 Thesis Contributions

This thesis develops a coherent sequence of decision-support models that progressively increase realism and policy relevance, moving from topology-driven network assessment to flow-aware, multi-modal, and fairness-aware optimization, and finally to cross-city benchmarking through safety-preserving network reduction. The unifying idea is to model and exploit network effects connectivity, routing, and demand feedback that create synergies in infrastructure design.

The first contribution is a safety-oriented graph representation of cycling infrastructure using a six-level Degree of Separation (DoS) classification aligned with a CycleRAP-style interpretation of safety and comfort, enabling reproducible network-scale analysis from widely available data. Building on this representation, the thesis identifies a small set of graph metrics that are both interpretable and suitable for optimization under upgrades, including a DoS-weighted notion of node strength, an arc-based closeness centrality defined on the line graph, and connectivity-based measures capturing whether high-DoS segments form a coherent backbone. These metrics are validated empirically against perception-based data and used to formulate a budget-constrained topology optimization framework that selects upgrades to improve network-level safety and comfort, with extensive evaluation across multiple French cities and comparisons to heuristic baselines.

The second contribution introduces cyclist demand into the planning problem through a flow-aware exposure model. Cyclist flows are estimated using an all-or-nothing assignment with a cyclist-specific generalized cost that captures attributes such as slope, turning angles, and traffic-related factors, and safety is measured via cyclist exposure as a flow-weighted sum of link-level exposure. The resulting bi-level structure captures how upgrades change route costs, reroute cyclists, and redistribute exposure, thereby quantifying network synergies driven by route-choice shifts and yielding upgrade portfolios that target corridors combining low DoS, high traffic exposure, and high predicted usage.

The third contribution extends the framework to a multi-modal setting by explicitly modeling car traffic assignment and congestion, allowing cycling upgrades to be evaluated not only for their safety benefits but also for their impact on car travel times under capacity changes and post-upgrade equilibrium flows. Results show that substantial safety improvements can be achieved with limited degradation of car performance when multi-modal interactions are internalized in the optimization.

The fourth contribution integrates induced modal shift and transport-justice objectives. By allowing some trips to switch from car to bicycle as cycling conditions improve, the model captures demand feedback that can partially offset congestion impacts, changing the assessment of which interventions are operationally acceptable. Two complementary optimization formulations are developed and compared: a utilitarian approach that improves aggregate system performance and a sufficientarian approach that prioritizes meeting minimum safety and accessibility thresholds. The sufficientarian formulation yields more balanced outcomes,

reducing safety and accessibility shortfalls while maintaining competitive overall efficiency.

Finally, the thesis addresses cross-city comparison by proposing a safety-preserving clustering framework that compresses large cycling networks into fixed-size clustered networks while preserving DoS, flows, and exposure-relevant attributes. Exposure-based kernels and distance measures computed on these reduced graphs provide label-invariant, size-robust similarity measures, enabling consistent benchmarking and monitoring of cycling-network safety structure across cities despite differences in network resolution and cardinality.

1.6 Publications

1.6.1 Journals

- Campero-Jurado, M., Canudas-de-Wit, C., & De Nunzio, G., "A Graph-Based Approach to the Topological Optimization of Cycling Networks for the Improvement of Safety and Comfort of Cyclists". *IEEE Transactions on Intelligent Transportation Systems (T-ITS).* (pp. 12558 - 12570), April, 2025, DOI: 10.1109/TITS.2025.3556555, Electronic ISSN: 1558-0016.

1.6.2 Working papers

- Campero-Jurado, M., Canudas-de-Wit, C., De Nunzio, G., & Salazar, M., "On Transport Justice in Bicycle Network Design: Balancing Safety and Car Impacts". *To be submitted in IEEE Transactions on Intelligent Transportation Systems (T-ITS).*
- Campero-Jurado, M., Canudas-De-Wit, C., De Nunzio, G.. "Clustering and Comparing Safety in Cycling Networks". *To be submitted in the IEEE 29th International Conference on Intelligent Transportation Systems (ITSC). IEEE*, Naples, Italy, September 2026.

1.6.3 International conferences

- Campero-Jurado, M., Canudas-De-Wit, C., De Nunzio, G., & Mourgues, R. (2024, September). "A general graph-based optimization framework for improving safety and comfort in cycling networks". *IEEE 27th International Conference on Intelligent Transportation Systems (ITSC) (pp. 124-129). IEEE*, Edmonton, Canada, September 2024.
- Campero-Jurado, M., Salazar, M., De Nunzio, G., & Canudas-de-Wit C. "Improving Urban Cycling Safety and Comfort through Optimized Infrastructure Upgrades". *23rd European Control Conference (ECC)*, Thessaloniki, Greece, June, 2025.
- Campero-Jurado, M., De Nunzio, G., & Canudas-De-Wit, C. (2025, November). "Safety-Oriented Design of Cycling Infrastructure Minimizing Impact on Vehicular Traffic". *28th International Conference on Intelligent Transportation Systems (ITSC).* Gold Coast, Australia, November, 2025.

- Campero-Jurado, M., Canudas-de-Wit, C., De Nunzio, G., & Salazar, M., "On Transport Justice and Safety in Bicycle Network Design Optimization", *IEEE Conference on Decision and Control (CDC)*, Rio de Janeiro, Brazil, December, 2025.

1.7 Thesis outline

This section provides the reader with an overview of the thesis structure and the role of each chapter. The thesis consists of the following chapters:

- Introduction
- Preliminaries
- Topological graph-based network design
- Network Design with Flow Redistribution
- Network Design with Flow Redistribution in a Multi-Modal Framework
- Network Design with Flow Redistribution in a Multi-Modal Network for Transport Justice
- Clustering and Comparing Safety in Cycling Networks Conclusions

Chapter 1 belongs to the first part and details the context and motivation, the state of the art, and the main contributions. The remainder of this section outlines each subsequent chapter, linking them to the contributions identified in Section 1.5 and to the publications listed in Section 1.6.

1.7.1 Chapter 2: Preliminaries

This chapter provides a concise reminder of the modeling components used throughout the thesis. Section 2.1 presents the weighted directed-graph representation adopted for the cycling and car networks and recalls the Degree of Separation (DoS) categorization used to encode safety and comfort; it also illustrates how DoS is distributed in the Grenoble cycling and car networks. Section 2.2 defines the cyclist generalized distance, which accounts for slope, angular change, and traffic-related factors. Section 2.3 describes the flow assignment procedures used to calculate cyclist and automobile flows, which support the metrics introduced in Section 2.4: cyclist exposure (derived from flows, link lengths, and traffic factors) and total car travel time. Finally, Section 2.5 summarizes the objective functions used across the different formulations considered in the thesis.

1.7.2 Chapter 3: Topological graph-based network design

This chapter develops a topology-driven framework to measure and improve cycling safety and comfort at network scale. Section 3.1.3 presents the perceived safety and comfort data used to ground and validate the modeling choices. Section 3.2 reviews graph-theory metrics commonly used to describe network structure and robustness, and highlights the non-trivial monotonicity requirement needed when arc weights represent safety levels and are upgraded. Section 3.3 introduces the framework used to predict perceived safety and comfort from the selected metrics. Section 3.3.3 formulates the budget-constrained network-upgrade optimization problem for improving DoS in order to maximize perceived safety and comfort. Section 3.4 presents the solution algorithms used throughout the thesis, in particular an ad hoc MLCC (maximize largest connected component) heuristic and a genetic algorithm warm-started from the MLCC solution (GA-U). Sections 3.5 and 3.5 compare the proposed methods with baseline heuristics on an academic benchmark, Sioux Falls network, and on the Grenoble cycling network, respectively. Section 3.5 extends the analysis to 14 French cities. Finally, Section 3.6 concludes and outlines future work.

The chapter is based on the following publication:

- Campero-Jurado, M., Canudas-de-Wit, C., & De Nunzio, G., "A Graph-Based Approach to the Topological Optimization of Cycling Networks for the Improvement of Safety and Comfort of Cyclists". *IEEE Transactions on Intelligent Transportation Systems (T-ITS).* (pp. 12558 - 12570), April, 2025, DOI: 10.1109/TITS.2025.3556555, Electronic ISSN: 1558-0016.
- Campero-Jurado, M., Canudas-De-Wit, C., De Nunzio, G., & Mourgues, R. (2024, September). "A general graph-based optimization framework for improving safety and comfort in cycling networks". *IEEE 27th International Conference on Intelligent Transportation Systems (ITSC) (pp. 124-129). IEEE*, Edmonton, Canada, September 2024.

1.7.3 Chapter 4: Network Design with Flow Redistribution

This chapter moves from purely topological planning to flow-aware decision support by incorporating cyclist demand and route choice. Section 4.2 introduces the cyclist generalized distance and the associated all-or-nothing (AoN) assignment used to estimate cyclist flows, and defines cyclist exposure as a product of flow, length, and traffic-related factors. Section 4.2.2 formulates, for the first time in the thesis, a bi-level optimization problem: the lower level computes flows induced by shortest-path routing under the cyclist distance, while the upper level selects DoS upgrades under a budget to minimize traffic exposure. Section 4.3 presents the Grenoble case study and evaluates the proposed modified genetic algorithm (GA-U) for exposure reduction. Section 4.4 concludes and discusses directions for future research.

The chapter is based on the following publication:

- Campero-Jurado, M., Salazar, M., De Nunzio, G., & Canudas-de-Wit C. "Improving Urban Cycling Safety and Comfort through Optimized Infrastructure Upgrades". *23rd European Control Conference (ECC)*, Thessaloniki, Greece, June, 2025.

1.7.4 Chapter 5: Network Design with Flow Redistribution in a Multi-Modal Framework

This chapter introduces a multi-modal perspective in which cycling upgrades may affect car performance through capacity reallocation. Section 5.2 presents the model design and the interaction mechanisms between cycling upgrades and car traffic within a bi-level formulation: the upper level reduces cyclist exposure while limiting deterioration in car travel times, and the lower level models cyclist route choice together with car-traffic responses under capacity constraints. Section 5.3 presents the solution algorithm for the bi-level problem applied to the Grenoble case study. Finally, Section 5.4 summarizes the main findings and outlines future work.

The chapter is based on the following publication:

- Campero-Jurado, M., De Nunzio, G., & Canudas-De-Wit, C. (2025, November). "Safety-Oriented Design of Cycling Infrastructure Minimizing Impact on Vehicular Traffic". *28th International Conference on Intelligent Transportation Systems (ITSC)*. Gold Coast, Australia, November, 2025.

1.7.5 Chapter 6: Network Design with Flow Redistribution in a Multi-Modal Network for Transport Justice

This chapter extends the multi-modal framework by incorporating induced modal shift and explicit transport-justice objectives. Section 6.2 introduces the mode-choice component through which improved cycling conditions can induce some drivers to switch to cycling, and describes the surrogate user equilibrium (UE) assignment used to handle demand changes efficiently. Section 6.2.3 presents a utilitarian formulation that jointly minimizes aggregate cyclist exposure and surrogate car travel time. Section 6.2.4 introduces a sufficientarian formulation that aims to increase the share of cyclists reaching a minimum safety threshold while limiting the share of drivers exceeding an accessibility (travel-time) threshold. Section 6.3 reports results on the Grenoble network and Section 6.4 concludes with final remarks.

The chapter is based on the following publication:

- Campero-Jurado, M., Canudas-de-Wit, C., De Nunzio, G., & Salazar, M., "On Transport Justice and Safety in Bicycle Network Design Optimization", *IEEE Conference on Decision and Control (CDC)*, Rio de Janeiro, Brazil, December, 2025.

- Campero-Jurado, M., Canudas-de-Wit, C., De Nunzio, G., & Salazar, M., "On Transport Justice in Bicycle Network Design: Balancing Safety and Car Impacts". ***To be submitted*** *in IEEE Transactions on Intelligent Transportation Systems (T-ITS).*

1.7.6 Chapter 7: Clustering and Comparing Safety in Cycling Networks

This chapter proposes a safety-preserving reduction framework to enable fair comparison of cycling networks across cities with different sizes and mapping resolutions. Section 7.2 presents the clustering framework used to build the reduced (supernetwork) representation. Section 7.3 introduces the algorithm used to cluster the larger network so that it matches the size of the smaller one. Section 7.4 introduces comparison metrics that remain meaningful under size differences and uses them to validate the clustering-based reduction. Section 7.5 illustrates the comparison framework on a toy example. Section 7.5 presents a case study on three French cities (Lyon, Strasbourg, and Grenoble). Finally, Section 7.6 concludes with remarks and perspectives.

The chapter is based on the following publication:

- Campero-Jurado, M., Canudas-De-Wit, C., De Nunzio, G.. "Clustering and Comparing Safety in Cycling Networks". ***To be submitted*** *in the IEEE 29th International Conference on Intelligent Transportation Systems (ITSC). IEEE*, Naples, Italy, September 2026.

1.7.7 Conclusions and future research

The thesis concludes with a summary of the work presented in the previous chapters and an outline of future research directions

Chapter 2

Preliminaries

This chapter introduces the notation and basic models used throughout the thesis. We first describe how transport networks are represented as weighted directed graphs, and how we distinguish the cycling layer from the car layer (including the possibility that the two layers share intersections and road segments). We then recall the *Degree of Separation* (DoS) classification, which provides an ordered six-level description of cycling facilities. Because all upgrade actions in this thesis are modeled as DoS increases on selected links, DoS serves as a common reference for infrastructure quality and for network improvement scenarios.

Next, we present the cyclist-demand setting and the flow model. Bicycle flows are estimated from origin–destination (OD) demand using an all-or-nothing (AoN) assignment on shortest paths computed with a cyclist-specific generalized distance (including length, slope, turning angles, and traffic-related factors). These flows are used to define a network-level cyclist-exposure measure, which captures where many cyclists ride and how exposed they are to motor traffic.

Finally, we introduce the car-traffic model used in the multi-modal chapters. Car travel times follow a congestion function with capacities, and car flows are computed under a user-equilibrium principle. Together, these definitions provide the shared modeling backbone for the chapters that follow, which progress from purely topological analysis to flow-aware, multi-modal, and fairness-oriented optimization.

2.1 Cycling Network Definition and Classification

We represent the transportation system as a *multi-modal* weighted directed graph (digraph) in order to capture both bicycle and automobile movements, and in particular to account for the fact that cycling-infrastructure upgrades may affect car traffic (e.g., through road-space reallocation). Formally, the full network is

$$\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{W}),$$

where $\mathcal{V}$ is the set of nodes (intersections, junctions, or road endpoints), $\mathcal{A}$ is the set of directed arcs (road segments), and $\mathcal{W}$ collects arc attributes/weights. The network is composed of two layers: a bicycle layer $\mathcal{G}^{\mathrm{b}}$ and a car layer $\mathcal{G}^{\mathrm{c}}$,

$$\mathcal{G}^{\mathrm{b}} = \big(\mathcal{V}^{\mathrm{b}}(\mathcal{G}^{\mathrm{b}}),\, \mathcal{A}^{\mathrm{b}}(\mathcal{G}^{\mathrm{b}}),\, \mathcal{W}^{\mathrm{b}}(\mathcal{G}^{\mathrm{b}})\big), \qquad \mathcal{G}^{\mathrm{c}} = \big(\mathcal{V}^{\mathrm{c}}(\mathcal{G}^{\mathrm{c}}),\, \mathcal{A}^{\mathrm{c}}(\mathcal{G}^{\mathrm{c}}),\, \mathcal{W}^{\mathrm{c}}(\mathcal{G}^{\mathrm{c}})\big).$$

The superscripts $(\cdot)^{\mathrm{b}}$ and $(\cdot)^{\mathrm{c}}$ indicate quantities associated with the bicycle and car layers, respectively. When no ambiguity arises, we simply write $\mathcal{V}^{\mathrm{b}}$, $\mathcal{A}^{\mathrm{b}}$, $\mathcal{W}^{\mathrm{b}}$ (and similarly for the car layer). The global graph satisfies

$$\mathcal{V} = \mathcal{V}^{\mathrm{b}} \cup \mathcal{V}^{\mathrm{c}}, \qquad \mathcal{A} = \mathcal{A}^{\mathrm{b}} \cup \mathcal{A}^{\mathrm{c}},$$

and the two layers may overlap:

$$\mathcal{V}^{\mathrm{b}} \cap \mathcal{V}^{\mathrm{c}} \neq \emptyset, \qquad \mathcal{A}^{\mathrm{b}} \cap \mathcal{A}^{\mathrm{c}} \neq \emptyset.$$

Bicycle layer. Nodes $\mathcal{V}^{\mathrm{b}}$ represent intersections and road endpoints, and arcs $\mathcal{A}^{\mathrm{b}}$ represent bike-lane segments. Each bicycle arc $a \in \mathcal{A}^{\mathrm{b}}$ is written as an ordered pair $a = (u, v)$ with $u, v \in \mathcal{V}^{\mathrm{b}}$ and carries (at least) two attributes: (i) a *degree-of-separation* (DoS) label $w_a^{\mathrm{b}} \in \{1, 2, 3, 4, 5, 6\}$, where larger values indicate stronger separation from motor traffic and thus safer and more comfortable cycling conditions; and (ii) a strictly positive geographical length $d_a^{\mathrm{b}} > 0$. Although $\mathcal{G}^{\mathrm{b}}$ is directed, we assume symmetric DoS labels whenever both directions exist, i.e., $w_{(u,v)}^{\mathrm{b}} = w_{(v,u)}^{\mathrm{b}}$. For any node $v \in \mathcal{V}^{\mathrm{b}}$, we denote by $\mathcal{N}_v \subseteq \mathcal{A}^{\mathrm{b}}$ the set of arcs incident to v. When referring to another bicycle graph (e.g., a subgraph $\mathcal{H}$), we explicitly write $\mathcal{V}^{\mathrm{b}}(\mathcal{H})$, $\mathcal{A}^{\mathrm{b}}(\mathcal{H})$, and $\mathcal{W}^{\mathrm{b}}(\mathcal{H})$.

Car layer. The car layer $\mathcal{G}^{\mathrm{c}}$ is defined analogously, with nodes $\mathcal{V}^{\mathrm{c}}$ and arcs $\mathcal{A}^{\mathrm{c}}$ representing the road network used by motor vehicles. Each car arc $a \in \mathcal{A}^{\mathrm{c}}$ is associated with a length $d_a^{\mathrm{c}} > 0$ and traffic-performance attributes (e.g., free-flow travel time and capacity) used later to compute congestion-dependent travel times. Unlike the bike layer, where weights capture safety and comfort, the car layer uses travel time as its generalized cost. Accordingly, when a generic arc-weight notation is needed, we set $w_a^{\mathrm{c}} = t_a^{\mathrm{c}}$, where t_a^{c} denotes the travel time required to traverse arc a.

Throughout the thesis, we use the generic notation a for an arc in either layer; when the mode matters, it is indicated by the superscripts b or c. This convention is particularly convenient for shared segments $a \in \mathcal{A}^{\mathrm{b}} \cap \mathcal{A}^{\mathrm{c}}$, which carry both bicycle (DoS-based) and car-traffic attributes. Finally, let $\mathcal{V}_{\mathcal{O}}^{\mathrm{b}}, \mathcal{V}_{\mathcal{D}}^{\mathrm{b}} \subseteq \mathcal{V}^{\mathrm{b}}$ denote the sets of origin and destination nodes for bicycle trips, and similarly let $\mathcal{V}_{\mathcal{O}}^{\mathrm{c}}, \mathcal{V}_{\mathcal{D}}^{\mathrm{c}} \subseteq \mathcal{V}^{\mathrm{c}}$ denote the sets of origins and destinations for car trips.

Why DoS is used throughout the thesis. Many classification frameworks quantify cycling comfort via traffic volume/speed, geometry, and separation (e.g., BSIR/RCI/IHS/B-CI/BLOS and the threshold-based LTS approach) [46–51]. In this work, DoS provides a practical separation-first abstraction that is transparent, easy to communicate, and portable at city scale. It is also compatible with CycleRAP-style safety/comfort scoring, allowing us to interpret w_a^{b} in a consistent way when comparing links, designing upgrade actions, and evaluating network-level outcomes. Fig. 2.1 visualizes the resulting categorization of Grenoble's cycling network and cars network in the same city.

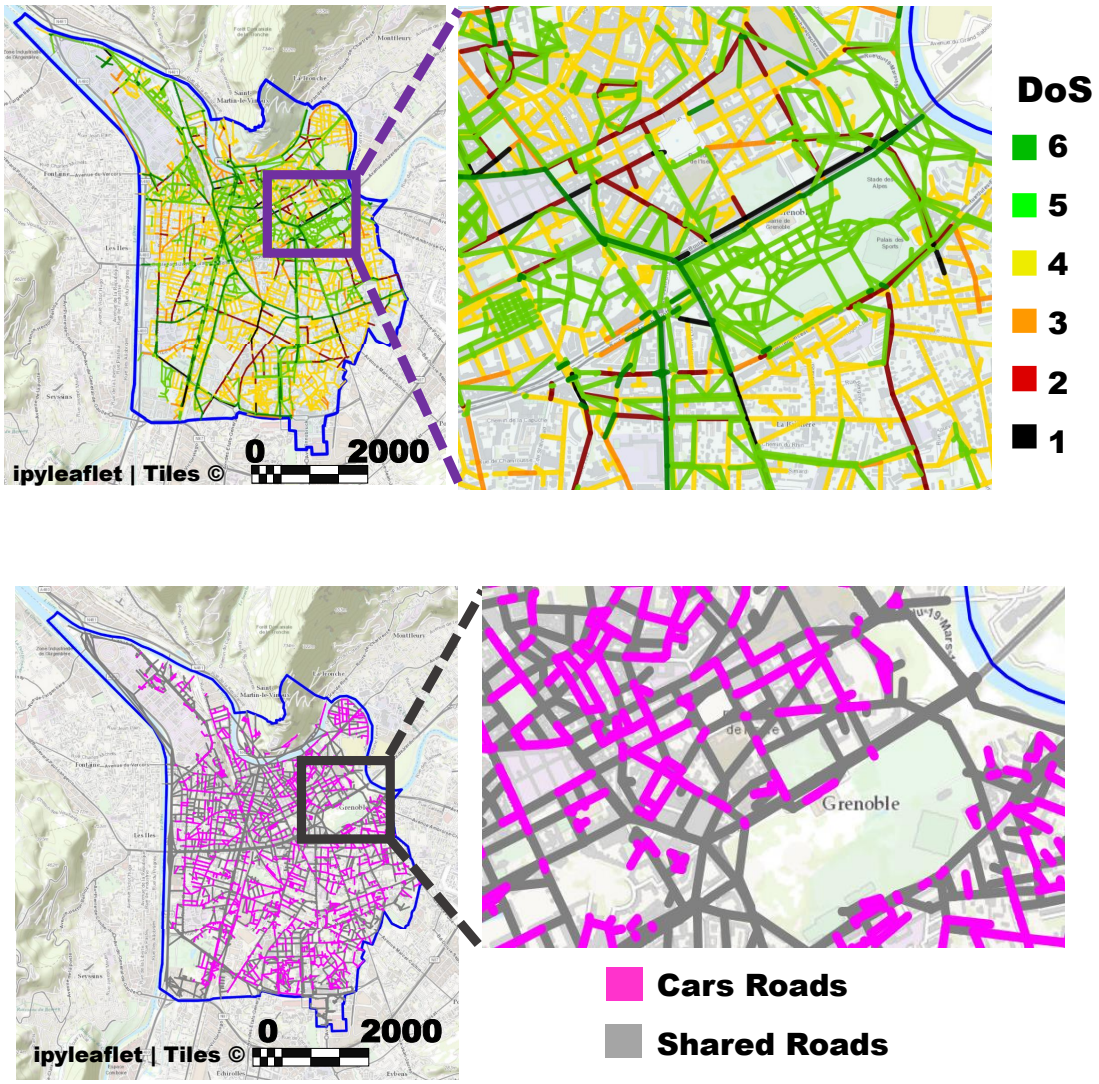

Figure 2.1: At the top: Cycling network of Grenoble, France, with the DoS ranging from 1 to 6, indicating increasing separation from motorized traffic. At the bottom: Road network for cars, including links shared with bicycles.

2.2 Cyclist Generalized Distance

To incorporate demand, we use the EMD2020 household travel survey [83] and define bicycle (OD) demand as follows. Let $\mathcal{V}^{\mathrm{b}}_{\mathcal{O}}$ and $\mathcal{V}^{\mathrm{b}}_{\mathcal{D}}$ be the sets of origin and destination nodes for bicycle trips. For each OD pair $(u, v) \in \mathcal{V}^{\mathrm{b}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{b}}_{\mathcal{D}}$, we denote the cycling demand by $q^{\mathrm{b}}_{(u,v)}$.

Cyclist routes are evaluated with a generalized distance that reflects cycling-relevant attributes beyond geometric length [84]. For a candidate path $p^{\mathrm{b}}_{(u,v)} = (a_1, \ldots, a_k)$ between u and v, each arc $a \in p^{\mathrm{b}}_{(u,v)}$ has length d^{b}_a, an angular change $\measuredangle^{\mathrm{b}}_a$ (along the arc and/or at

junctions), a slope factor s_a^{b}, and a traffic factor f_a^{b} capturing the effect of adjacent motor traffic and separation. The generalized cyclist distance from u to v is

$$\bar{d}^{\mathrm{b}}_{(u,v)} = \sum_{a \in p^{\mathrm{b}}_{(u,v)}} \left(d_a^{\mathrm{b}}\, s_a^{\mathrm{b}}\, (f_a^{\mathrm{b}})^{\rho} \right) + \sum_{a \in p^{\mathrm{b}}_{(u,v)}} \measuredangle_a^{\mathrm{b}}\, \kappa, \tag{2.1}$$

where $\rho > 0$ and $\kappa > 0$ are calibration parameters controlling the influence of traffic exposure and angular change, respectively. For convenience, the per-arc contribution is

$$\bar{d}_a^{\mathrm{b}} = d_a^{\mathrm{b}} \cdot s_a^{\mathrm{b}} \cdot (f_a^{\mathrm{b}})^{\rho} + \measuredangle_a^{\mathrm{b}} \cdot \kappa. \tag{2.2}$$

The traffic factor depends on the DoS label and on adjacent traffic intensity (via AADT) when separation is incomplete; for highly separated facilities (DoS 5–6), AADT effects are neutralized [5]:

$$f_a^{\mathrm{b}} = \begin{cases} \psi\, \exp(\mathrm{AADT}(a)), & \text{if } w_a^{\mathrm{b}} < 5, \\ \psi \cdot 2.03, & \text{if } w_a^{\mathrm{b}} = 5, \\ \psi, & \text{if } w_a^{\mathrm{b}} = 6, \end{cases} \tag{2.3}$$

where ψ is a scaling factor and the factor 2.03 reflects the increased crash risk in pedestrian-shared spaces compared with fully separated facilities [85].

2.3 Flow Assignment in Bicycle and Car Layers

2.3.1 Bicycle All-or-Nothing Assignment

Following [56], we approximate cyclist route choice using an all-or-nothing (AoN) assignment, which is appropriate when congestion effects on bicycle travel times are negligible. Under AoN, each OD demand $q^{\mathrm{b}}_{(u,v)}$ is assigned to a single shortest path $p^{\mathrm{b}}_{(u,v)}$ computed with respect to $\bar{d}^{\mathrm{b}}_{(u,v)}$ in Eq. (2.1).

Let $\mathbb{I}_{a \in p^{\mathrm{b}}_{(u,v)}}$ be an indicator equal to 1 if arc a lies on the chosen shortest path between u and v, and 0 otherwise. The resulting cyclist flow on arc a is

$$\phi_a^{\mathrm{b}} = \sum_{(u,v) \in \mathcal{V}^{\mathrm{b}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{b}}_{\mathcal{D}}} q^{\mathrm{b}}_{(u,v)} \cdot \mathbb{I}_{a \in p^{\mathrm{b}}_{(u,v)}} \;\geq\; 0. \tag{2.4}$$

2.3.2 Car User-Equilibrium Assignment

As mentioned in Chapter 2.1, for car trips, $\mathcal{V}^{\mathrm{c}}_{\mathcal{O}}$ and $\mathcal{V}^{\mathrm{c}}_{\mathcal{D}}$ denote OD sets, and let $q^{\mathrm{c}}_{(u,v)}$ denote car demand between $(u,v) \in \mathcal{V}^{\mathrm{c}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{c}}_{\mathcal{D}}$. Unlike the bicycle layer, car travel times depend on congestion. We model car route choice through a user-equilibrium (UE) assignment [86], where no driver can reduce travel time by unilaterally changing routes.

Each car arc $a \in \mathcal{A}^{\mathrm{c}}$ has flow ϕ_a^{c} and travel time $t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}})$. We use the Bureau of Public Roads (BPR) function [87]

$$t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}}) = t_a^0 \left(1 + \mu \left(\frac{\phi_a^{\mathrm{c}}}{\theta_a^{\mathrm{c}}}\right)^{\sigma}\right), \tag{2.5}$$

where t_a^0 is the free-flow travel time, θ_a^{c} is the arc capacity, and μ, σ are standard parameters. UE flows can be computed with standard traffic-assignment solvers based on this cost function; the result is a consistent set of link flows $\{\phi_a^{\mathrm{c}}\}_{a\in\mathcal{A}^{\mathrm{c}}}$ and link travel times $\{t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}})\}_{a\in\mathcal{A}^{\mathrm{c}}}$.

2.4 Safety and Efficiency Metrics

2.4.1 Cyclist Exposure

To measure cycling safety in a way that accounts for who is affected, we use cyclist exposure. At the network scale, exposure is the flow-weighted sum of arc exposures:

$$\mathcal{E} = \sum_{a\in\mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}} \cdot f_a^{\mathrm{b}} \cdot d_a^{\mathrm{b}}, \tag{2.6}$$

so exposure increases on links that (i) carry more cyclists, (ii) are more exposed to motor traffic (via f_a^{b}), and/or (iii) are longer.

2.4.2 Total Car travel Time

To quantify motor-traffic performance, we use total car travel time:

$$\mathcal{TTT}^{\mathrm{c}} = \sum_{a\in\mathcal{A}^{\mathrm{c}}} \phi_a^{\mathrm{c}} \cdot t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}}), \tag{2.7}$$

where $t_a^{\mathrm{c}}(\cdot)$ is given by Eq. (2.5). This metric aggregates congestion impacts over the full car network and is used in later chapters to evaluate the trade-off between cycling-safety benefits (via $\mathcal{E}$) and motor-traffic impacts (via $\mathcal{TTT}^{\mathrm{c}}$) under capacity changes induced by cycling upgrades.

2.5 Objective Functions

Across the thesis, upgrade decisions are modeled as discrete increases of DoS labels on bicycle arcs. Let $\boldsymbol{w}^{\mathrm{b}} = [w_a^{\mathrm{b}}]_{a\in\mathcal{A}^{\mathrm{b}}}$ denote the original DoS assignment on the bicycle network $\mathcal{G}^{\mathrm{b}}$, and let $\tilde{\boldsymbol{w}}^{\mathrm{b}} = [\tilde{w}_a^{\mathrm{b}}]_{a\in\mathcal{A}^{\mathrm{b}}}$ be a candidate upgraded configuration. Feasible upgrade configurations are collected in the set

$$\Omega = \left\{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \{1,\dots,6\}^{|\mathcal{A}^{\mathrm{b}}|} : w_a^{\mathrm{b}} \le \tilde{w}_a^{\mathrm{b}} \le 6,\ \forall a \in \mathcal{A}^{\mathrm{b}}\right\},$$

that is, each arc can only keep its current DoS level or be upgraded to a higher one, up to the maximum level 6. Throughout the chapters, the planner selects $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega$ under a global budget constraint and optimizes a network-level objective that formalizes the intended notion of "improvement". For clarity, we summarize below the objective functions used in each formulation.

2.5.1 Perception-Based Improvement (Chapter 3)

In Chapter 3, the objective is to improve the predicted (perceived) safety and comfort provided by the regression model. Let $\hat{y}_{\boldsymbol{w_0}}$ be the predicted output for the original configuration $\boldsymbol{w_0}$ and $\hat{y}(\tilde{\boldsymbol{w}})$ the predicted output under upgrades $\tilde{\boldsymbol{w}}$. The upgrade objective is defined as the predicted gain in perceived safety and comfort:

$$\|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\|\,. \tag{2.8}$$

The network-upgrade problem in this chapter is therefore posed as

$$\tilde{\boldsymbol{w}}^{\mathrm{b}\star} = \arg\max_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega} \|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\| \qquad \text{s.t.} \quad \sum_{a \in \mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \leq B. \tag{2.9}$$

2.5.2 Exposure Minimization (Chapter 4)

In Chapter 4, the objective moves beyond topology and adopts a flow-aware safety criterion. For any candidate upgrade configuration, bicycle flows ϕ_a^{b} are computed using an all-or-nothing assignment based on the cyclist generalized distance. The safety objective is the total cyclist exposure:

$$\mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}} \cdot f_a^{\mathrm{b}} \cdot d_a^{\mathrm{b}}. \tag{2.10}$$

The bi-level formulation minimizes $\mathcal{E}$ at the upper level, subject to the budget constraint.

2.5.3 Exposure–Travel-Time Trade-Off (Chapter 5)

In Chapter 5, cycling upgrades may reduce road capacity on shared segments, which can increase car travel times. The objective therefore combines (i) cyclist exposure and (ii) a surrogate for total car travel time. Using a trade-off parameter $\gamma \in [0, 1]$, the objective function is:

$$\gamma \cdot \mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + (1 - \gamma) \cdot \mathcal{TTT}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}), \tag{2.11}$$

where $\mathcal{TTT}^{\mathrm{c}}$ is computed from car flows (fixed in this chapter) and capacity updates induced by high-DoS upgrades on shared arcs. The optimization minimizes this objective under the budget constraint.

2.5.4 Utilitarian Objective (Chapter 6)

In Chapter 6, demand can respond through modal shift, and car performance is evaluated using a network-wide travel-time indicator. The utilitarian objective retains the same weighted sum of exposure and travel time, but here $\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ denotes a predicted total travel time obtained from a graph neural network model that accounts for modal shift from cars to bikes:

$$\gamma \cdot \mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + (1-\gamma) \cdot \mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}). \tag{2.12}$$

This objective reflects an average performance perspective: it reduces cyclist exposure while limiting aggregate losses in car and bicycle efficiency.

2.5.5 Sufficientarian Objective (Chapter 6)

To explicitly address distributional concerns, Chapter 6 also introduces a sufficientarian formulation based on insufficiency indicators. The objective combines (i) safety insufficiency and (ii) time-accessibility insufficiency for both cyclists and drivers:

$$\gamma \cdot J_{\mathcal{SI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + (1-\gamma) \cdot \Big(J^{\mathrm{b}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + J^{\mathrm{c}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \Big). \tag{2.13}$$

Here, $J_{\mathcal{SI}}$ penalizes the share of cyclists whose feasible routes remain below a safety threshold, while $J^{\mathrm{b}}_{\mathcal{TI}}$ and $J^{\mathrm{c}}_{\mathcal{TI}}$ penalize the share of cyclists and drivers, respectively, whose travel times exceed an accessibility threshold. The optimization minimizes this objective under the same budget constraint and capacity-update rules.

Remark. In all chapters, the budget constraint has the common form

$$\sum_{a \in \mathcal{A}^{\mathrm{b}}} d^{\mathrm{b}}_a \cdot \pi(w^{\mathrm{b}}_a, \tilde{w}^{\mathrm{b}}_a) \leq B, \tag{2.14}$$

and the decision variables satisfy $w^{\mathrm{b}}_a \leq \tilde{w}^{\mathrm{b}}_a \leq 6$ for all $a \in \mathcal{A}^{\mathrm{b}}$.

Chapter 3

Topological graph-based network design

Contents

3.1 Introduction . . . 35
3.1.1 Summary of Results . . . 36
3.1.2 Related Work . . . 37
3.1.3 Perceived Safety and Comfort Data . . . 38
3.2 Centrality Measures for Safety and Comfort . . . 39
3.2.1 Network Centrality Measures . . . 40
3.2.2 Monotonicity under DoS Upgrades . . . 43
3.2.3 Numerical Monotonicity Test . . . 45
3.2.4 Final Metrics for Predicting Perceived Safety and Comfort . . . 46
3.3 Model Design . . . 49
3.3.1 Linear Regression Approach . . . 50
3.3.2 Model complexity and evaluation criteria. . . . 53
3.3.3 Linear Regression vs. Ensemble ML Models . . . 54
3.4 Solution Algorithms . . . 57
3.5 Benchmark Instances and Case Studies . . . 61
3.6 Concluding Remarks . . . 68

3.1 Introduction

Cycling offers well-established health, environmental, and socio-economic benefits, including improved public health outcomes and reduced emissions from short car trips [1]. Nevertheless, in many cities its everyday use remains limited, largely because potential riders perceive cycling as unsafe or uncomfortable in mixed-traffic environments [2]. A central determinant of both perceived and objective safety is the degree of separation between cyclists and motorized traffic: facilities that provide stronger physical separation are consistently associated with higher comfort and greater willingness to cycle [3, 4].

Because upgrading cycling infrastructure is costly and budgets are constrained, cities need decision-support tools that can (i) quantify safety and comfort in a way that is consistent across a network and (ii) identify which upgrades deliver the greatest benefits per unit of investment. Graph-based representations of transport networks are natural for this purpose: they capture connectivity, redundancy, and the availability of low-stress alternatives, and they enable the use of established metrics from complex network analysis [6]. However, when safety is encoded as arc weights (here, separation levels), an additional requirement becomes critical for optimization: the chosen indicators should behave predictably under upgrades, ideally, improving or at least not degrading when an arc is made safer. This motivates a careful study of the monotonicity of candidate metrics under DoS increases, which is fundamental when upgrades are modeled as weight improvements in a transport network.

In this chapter, we first review a set of graph-theory metrics commonly used to characterize network structure and robustness [6]. We then investigate their monotonic behavior under random safety upgrades in a weighted graph, highlighting which indicators are suitable for an optimization setting. Based on this analysis, we select the metrics that best capture cycling safety and comfort when arc weights represent the degree of separation, and we place them in the context of the CycleRAP assessment methodology [5]. Finally, we formulate an optimization framework to identify bike-lane upgrades that most effectively improve network-level safety and comfort under budget constraints, and we compare several algorithmic strategies to solve the problem in practice.

3.1.1 Summary of Results

This work develops and validates a graph-based framework to assess and improve cycling-network safety and comfort under budget constraints. Using OpenStreetMap data, the cycling infrastructure is first represented as a weighted graph and classified with a six-level DoS category. These DoS labels are then scored with CycleRAP to obtain a consistent, safety-oriented weighting of links. A main result is the identification of graph indicators that best reflect cyclists' perceived safety and comfort. Several candidates are evaluated and compared, and the study retains a subset of metrics, most notably node strength, an arc-based closeness centrality (in the line-graph formulation), and high-DoS path connectivity measures such as the largest connected component, because they are the most informative and usable for optimization. The chapter then establishes an empirical link between these network metrics and perceived safety/comfort, by aligning them with collected survey data. This yields an explicit model that captures real perceptions and can be used as an objective function for planning. Building on this model, the work proposes a topology optimization algorithm that selects bike-lane upgrades to maximize perceived safety and comfort within a limited budget, and is designed to be transferable to other cities. The approach is validated through comparisons with well-known heuristic solvers, including genetic algorithms and Bayesian optimization, and (when feasible) against a brute-force baseline. Across real case studies applied to 14 French cities, the results show that the proposed method consistently improves perceived safety and comfort while maintaining a strong trade-off between solution quality and computational cost, making it practical for decision support in urban planning.

3.1.2 Related Work

Our work relates to three strands of research: (i) the quantification of cycling safety and comfort, (ii) the optimization of cycling infrastructure under limited resources, and (iii) the use of graph-theory metrics to characterize and improve cycling networks.

Safety and comfort modeling: A substantial body of work proposes indices and classification schemes to rate cycling conditions. The Levels of Traffic Stress (LTS) framework classifies streets into discrete stress levels based on roadway design and traffic conditions [88]. Complementary approaches such as the Bicycle Safety Index Rating (BSIR) and the Bicycle Level of Service (BLOS) incorporate perceived safety, comfort, and operational conditions [89, 90]. Although widely used, these models often depend on local calibration, survey instruments, or context-specific data, which can limit the transferability between cities and countries. To address this limitation, we adopt a DoS classification that maps cycling facilities into six categories purely from OpenStreetMap (OSM) attributes, making it readily applicable across regions [54]. In this representation, higher DoS corresponds to greater physical separation from motorized traffic and, consequently, higher expected safety and comfort.

Cycling network optimization: Many studies investigate how to allocate infrastructure investments to improve cycling conditions while satisfying budget constraints. The Bicycle Network Improvement Problem (BNIP) and related formulations typically aim to enhance safety, reduce detours, and increase network usability through targeted upgrades [91, 92]. Safety and compatibility are commonly operationalized via stress/comfort indices such as LTS and the Bicycle Compatibility Index (BCI) [88, 90]. Other objective functions emphasize maximizing cycling facility usage [56, 57], improving connectivity and accessibility, or optimizing benefit–cost trade-offs [60, 61]. Demand-driven models further integrate cycling and motor-vehicle networks to guide where improvements may yield the greatest behavioral impact [65]. Methodologically, mixed-integer optimization, metaheuristics, and evolutionary approaches have all been explored to balance user preferences, travel time, and investment costs [59, 93, 94]. Despite this progress, many formulations rely on proxies for safety/comfort that are not consistently validated against transferable infrastructure-based measures, motivating our use of a DoS-weighted network representation.

Graph-theory network metrics: Graph metrics provide compact descriptors of network structure and have been used to assess cycling network performance and guide interventions. Betweenness centrality highlights critical nodes or links on shortest paths [95]; algebraic connectivity and spectral quantities capture robustness and efficiency [96, 97]; and modularity-based analyses identify community structure that can support targeted investments [96]. These tools have been applied to study accessibility, resilience, and the interaction of bicycle and vehicle systems, including clustering-based design strategies [98–102]. However, improvements in infrastructure do not necessarily translate monotonically into better graph metrics, which may respond nonlinearly to local modifications. This motivates optimization formulations that explicitly model safety/comfort attributes on the network, rather than assuming that generic topological improvements always align with cycling-level outcomes.

3.1.3 Perceived Safety and Comfort Data

We collect perceived safety and comfort data from the "Baromètre Vélo" (Bicycle Barometer) project [103], a satisfaction survey designed to assess and promote bicycle use in various cities and regions of France. This questionnaire, conducted every two years, collects cyclist perceptions of their cycling environment. The survey, open to all frequent visitors of a municipality, includes around 30 questions on general impressions, safety, comfort, cycling importance, and parking. Responses are geo-referenced to identify high-risk areas, with ratings reflecting perceived safety and infrastructure quality. The survey generates global scores for municipalities, including perceived safety and comfort. Fig. 3.1 illustrates the "Baromètre Vélo" for Grenoble, the global scores from 3009 survey respondents are shown on the left. The red, green and blue dots also originate from the survey responses: Red dots indicate priority areas for improvement, green mark perceived improvements over the past two years, and blue represent desired bike parking. While this work concentrates on global scores, future work will delve into the local insights offered by the survey, as represented by these colored dots.

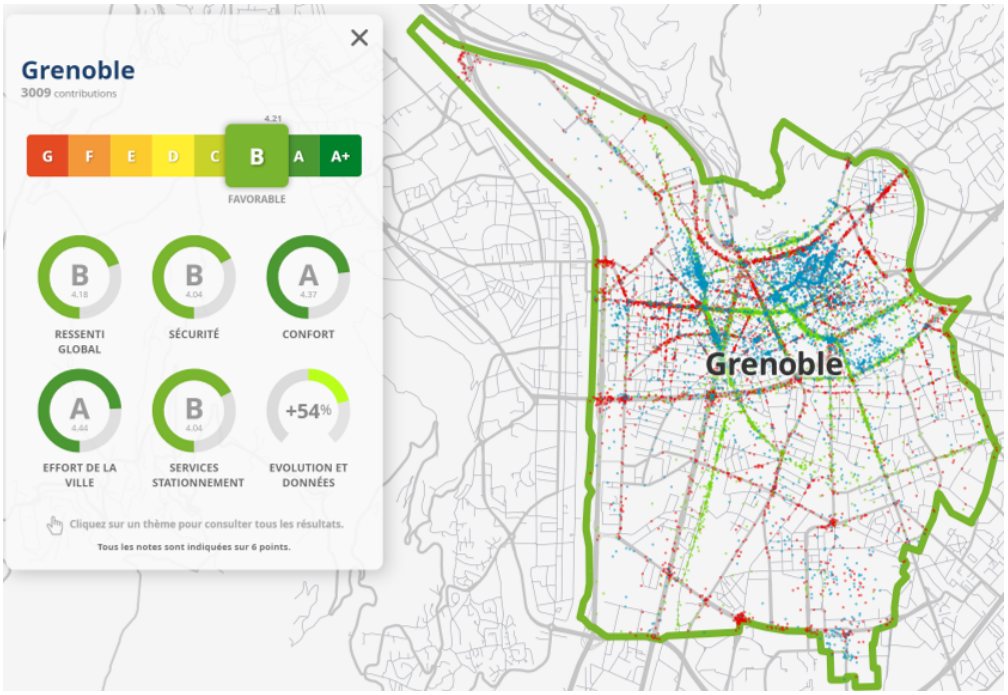

Figure 3.1: Perceived cycling safety and comfort data in Grenoble, France.

Having introduced the perceived-based safety and comfort indicators that serve as our reference signal, we now turn to graph-theory metrics. In the following section, we review a set of candidate network metrics, identify those that satisfy the required monotonicity properties under DoS upgrades, and explain why they are relevant for capturing cycling safety and comfort at the city scale. We then evaluate these candidates using the Baromètre Vélo scores and retain the most suitable ones for the optimization framework developed in this chapter.

3.2 Centrality Measures for Safety and Comfort

Characterizing the structural properties of a cycling network is essential for understanding how its topology may influence perceived safety and comfort. Graph-theory metrics provide a well-founded way to quantify positions of importance, structural cohesion, bottlenecks, or accessibility within a network, and an extensive literature has shown their usefulness for describing transport systems.

In the context of transport networks, such metrics have been used to capture a wide range of structural properties. Betweenness centrality identifies key nodes along shortest paths, highlighting critical infrastructure points where flows tend to concentrate [95]. Algebraic connectivity, measured by the second smallest Laplacian eigenvalue, is employed to assess robustness and vulnerability to disruptions [96], while the spectral radius is used to evaluate communication and propagation efficiency across the network [97]. Modularity helps detect clusters or communities of nodes, supporting the identification of sub-networks that may require targeted interventions [96]. Together, these metrics quantify connectivity, efficiency, and resilience, and have been applied to guide network improvement strategies such as prioritizing upgrades on critical paths, minimizing vulnerabilities, and promoting equitable access across diverse regions [98–100]. Graph-theory approaches have also been used more specifically to study bike–vehicle interactions and to develop cycling network design strategies based on clustering and corridor identification [101, 102].

However, most of these approaches implicitly assume that changes to the cycling infrastructure, for example, upgrading bike lanes or improving connectivity, will consistently improve the chosen graph metrics. This overlooks the fact that many centrality and spectral measures can exhibit complex, non-local, and non-monotonic behavior under local modifications of the network. As a result, it is not guaranteed that making parts of the network safer or more comfortable (in terms of DoS) will translate into a coherent improvement in the corresponding graph-theory metric.

Yet, our ultimate goal is to identify metrics that relate closely to how cyclists experience urban space. Since safety and comfort perceptions emerge from the interplay between infrastructure design, connectivity, exposure to traffic, and available route alternatives, our first task is therefore to evaluate a set of widely used graph-theory metrics and examine whether they can both (i) meaningfully predict observed safety and comfort ratings and (ii) respond monotonically to DoS upgrades. This motivates the systematic study of different centrality measures presented in the following sections.

Following the overview provided in *Network centrality: an introduction* [6], we focus on six commonly applied metrics, degree, k-core, closeness, betweenness, eigenvector, and accessibility centrality, each reflecting a distinct topological characteristic of the network. We first recall their definitions using a consistent notation, then examine their behavior under DoS upgrades. This monotonicity analysis will motivate the three centrality-based indicators ultimately employed in our optimization framework: node strength, arc-based closeness on the line graph, and a path-level metric based on largest connected components.

DoS : CSCS	Description	AADT	Mean speed
6: 5.23	Cycleways exclusively reserved for bikes, physically separated from cars.	14 400	50 km/h
5: 4.64	Cycleways separated from motor vehicles but shared with pedestrians.	16 000	50 km/h
4: 2.99	Roads where cycling is allowed, no designated lanes, shared with cars (speed limit: 30 km/h) .	960	30 km/h
3: 2.18	Shared roads allow cycling alongside motor vehicles, with cycling pictograms.	3120	30 km/h
2: 1.55	Roads with painted bike lanes.	6 912	40 km/h
1: 1.00	Unmarked shared car-bike lanes.	8 448	40 km/h

Table 3.1: The table displays the degree of separation (DoS) with its CycleRAP Safety and Comfort Score (CSCS), cycling infrastructure category, and GTLVille data on AADT and mean speed for roads near cycling paths in Grenoble.

3.2.1 Network Centrality Measures

As introduced in the preliminaries in Chapter 2, we represent the cycling infrastructure by a weighted directed graph

$$\mathcal{G}^{\mathrm{b}} = \big(\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}}\big), \tag{3.1}$$

where $\mathcal{V}^{\mathrm{b}}$ is the set of nodes (intersections/road endpoints) and $\mathcal{A}^{\mathrm{b}}$ the set of directed arcs (bike-lane segments). Each arc $a = (u, v) \in \mathcal{A}^{\mathrm{b}}$ is associated with a DoS label $w_a^{\mathrm{b}} \in \{1, \ldots, 6\}$ (larger values indicate greater separation from motor traffic) and a geographical length $d_a^{\mathrm{b}} > 0$. To incorporate a richer safety interpretation than the raw DoS scale, we map DoS levels to CycleRAP Safety and Comfort Scores (CSCS) through a strictly increasing function

$$h : \{1, \ldots, 6\} \to \mathbb{R}_{>0}, \tag{3.2}$$

whose values are reported Table 3.1, and Fig. 3.2 shows an example of how CycleRAP platform works. The mapping is obtained by inverting the original CycleRAP risk scores (so that larger values correspond to safer and more comfortable conditions) and normalizing them such that $h(1) = 1$. Hence, each arc a is assigned the weight $h(w_a^{\mathrm{b}})$, which serves as the safety-oriented weight used in the network metrics and optimization models developed in this thesis.

For metrics that require a notion of distance or cost, such as closeness and betweenness centrality, we construct an effective arc cost $\ell(a)$ that decreases as the DoS (and hence the CSCS value $h(w_a^{\mathrm{b}})$) increases. A simple and convenient choice is

$$\ell(a) = d_a^{\mathrm{b}} \cdot g(w_a^{\mathrm{b}}), \tag{3.3}$$

where $g(\cdot)$ is a monotonically decreasing function of the DoS level. One choice, consistent with the CSCS mapping, is

$$g(w_a^{\mathrm{b}}) = h(7 - w_a^{\mathrm{b}}), \tag{3.4}$$

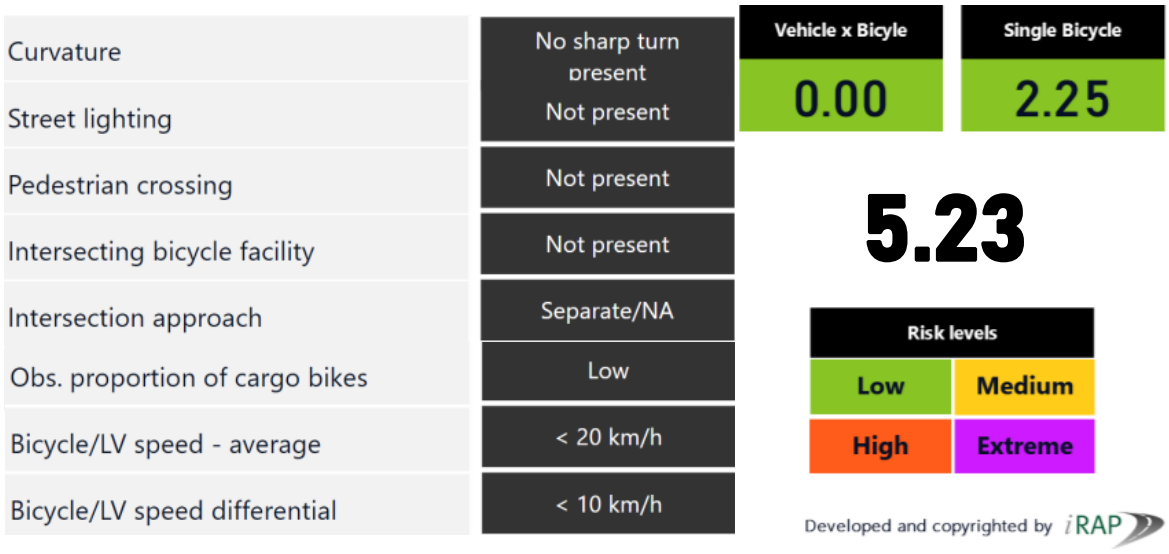

Figure 3.2: CycleRAP demonstrator tool for scoring cycling facilities.

which assigns to DoS level w_a^{b} the CSCS corresponding to the complementary level $7 - w_a^{\mathrm{b}}$ and is therefore strictly decreasing in w_a^{b}. With this construction, safer high-DoS links correspond to shorter effective distances, so that shortest-path computations naturally favor low-unsafety routes. We now introduce several graph-theory metrics commonly used to analyze the topological characteristics of complex networks [6].

i) Weighted degree centrality: The (unweighted) degree centrality of $v \in \mathcal{V}^{\mathrm{b}}$ counts the number of incident arcs:

$$c_{\deg}(v) = |\mathcal{N}_v|. \tag{3.5}$$

In a weighted setting with DoS labels w_a^{b}, it is natural to consider the strength of v:

$$c_{\mathrm{str}}(v) = \sum_{a \in \mathcal{N}_v} h(w_a^{\mathrm{b}}), \tag{3.6}$$

which aggregates the quality of the links meeting at v. Nodes with high $c_{\mathrm{str}}(v)$ are connected to many high-DoS links, and thus correspond to intersections embedded within well-protected parts of the network.

ii) k-core (weighted coreness): In this work, the k-core of $\mathcal{G}^{\mathrm{b}}$ is defined in a weighted sense by using node strength. For any integer threshold $k \geq 0$, the k-core of $\mathcal{G}^{\mathrm{b}}$, denoted $\mathcal{G}_k^{\mathrm{b}}$, is the maximal induced subgraph such that every node $v \in \mathcal{V}^{\mathrm{b}}$ satisfies $c_{\mathrm{str}}(v) \geq k$. The (weighted) coreness of a node v is then

$$c_k(v) = \max\{\, k : \; v \in \mathcal{G}_k^{\mathrm{b}} \,\}. \tag{3.7}$$

This metric highlights structurally cohesive regions supported by large arc weights (e.g., highly separated cycling infrastructure).

iii) Weighted closeness centrality: Weighted closeness centrality measures how close a node is to all others when distances along the network incorporate the effective costs $\ell(a)$

associated with safety and comfort. Let $d_\ell(v,u)$ denote the length of a shortest path from v to u with respect to arc costs $\ell(a)$. The closeness of v is then

$$c_{\text{clo}}(v) = \frac{|\mathcal{V}^{\text{b}}| - 1}{\sum_{u \in \mathcal{V}^{\text{b}} \setminus \{v\}} d_\ell(v,u)}. \tag{3.8}$$

Nodes with high $c_{\text{clo}}(v)$ enjoy short short low-risk routes (i.e., high-DoS/high-CSCS) routes to the rest of the network, and thus provide efficient access to safe corridors.

iv) Weighted betweenness centrality: Weighted betweenness centrality evaluates how often a node lies on shortest paths when path lengths use the DoS-aware costs $\ell(a)$. Let σ_{st}^{ℓ} be the number of shortest paths between nodes s and t with respect to $\ell(a)$, and let $\sigma_{st}^{\ell}(v)$ be the number of such paths that pass through v. The betweenness of v is

$$c_{\text{bet}}(v) = \sum_{\substack{s \neq v \neq t \\ s,t \in \mathcal{V}^{\text{b}}}} \frac{\sigma_{st}^{\ell}(v)}{\sigma_{st}^{\ell}}. \tag{3.9}$$

Nodes with high $c_{\text{bet}}(v)$ act as critical connectors in the weighted sense, concentrating many low-cost (high-DoS/high-CSCS) routes or bridging otherwise weakly connected regions.

v) Weighted eigenvector centrality: Let W be the weighted adjacency matrix of $\mathcal{G}^{\text{b}}$, with

$$W_{uv} = \begin{cases} h(w_a^{\text{b}}), & \text{if } a = (u,v) \in \mathcal{A}^{\text{b}}, \\ 0, & \text{otherwise.} \end{cases} \tag{3.10}$$

The weighted eigenvector centrality is the vector $c_{\text{eig}} \in \mathbb{R}^{|\mathcal{V}^{\text{b}}|}$ satisfying

$$W\, c_{\text{eig}} = \lambda_{\max} \cdot c_{\text{eig}}, \tag{3.11}$$

where $\lambda_{\max}$ is the largest eigenvalue of W. Equivalently, for each node $v \in \mathcal{V}^{\text{b}}$,

$$c_{\text{eig}}(v) = \frac{1}{\lambda_{\max}} \sum_{u \in \mathcal{V}^{\text{b}}} W_{vu} \cdot c_{\text{eig}}(u) = \frac{1}{\lambda_{\max}} \sum_{(v,u) \in \mathcal{A}^{\text{b}}} h(w_{(v,u)}^{\text{b}}) \cdot c_{\text{eig}}(u). \tag{3.12}$$

Nodes with high $c_{\text{eig}}(v)$ are strongly connected, directly or indirectly, to other nodes via high-DoS arcs, indicating influence within globally safe regions of the network.

vi) Weighted accessibility centrality: Accessibility centrality measures how diversely a node can reach others through random walks biased by arc weights. Define the transition probabilities of a one-step random walk as

$$\mathbb{P}(v \to u) = \frac{h(w_{(v,u)}^{\text{b}})}{\sum_{(v,x) \in \mathcal{A}^{\text{b}}} h(w_{(v,x)}^{\text{b}})}, \tag{3.13}$$

so that higher-DoS (higher-CSCS) arcs are more likely to be taken. Let $\mathbb{P}_v^{(t)}(u)$ denote the probability that a random walk starting from v is at node u after t steps. The accessibility of v (at horizon t) is defined as the exponential of the Shannon entropy of this distribution:

$$c_{\text{acc}}^{(t)}(v) = \exp\left(- \sum_{u \in \mathcal{V}^{\text{b}}} \mathbb{P}_v^{(t)}(u)\, \ln \mathbb{P}_v^{(t)}(u) \right). \tag{3.14}$$

High $c_{\mathrm{acc}}^{(t)}(v)$ values indicate that v can diffuse flow into the network through many balanced, high-DoS directions, reflecting intuitive, low-stress route availability.

Together, these six centralities provide a multifaceted characterization of node importance, from local connectivity (degree/strength) and structural embedding (k-core) to global influence (betweenness, eigenvector), efficiency (closeness), and dynamic reachability (accessibility). In Fig. 3.3 we illustrate these centralities with an academic example in which the DoS of the links were assigned randomly.

3.2.2 Monotonicity under DoS Upgrades

Our objective is not merely to characterize the topology of the cycling network, but to employ graph-theory metrics as optimization objectives in a setting where arcs are upgraded between different DoS categories. Recall that each arc $a \in \mathcal{A}^{\mathrm{b}}$ is assigned a DoS level $w_a^{\mathrm{b}} \in \{1, \ldots, 6\}$, with higher values indicating safer, more comfortable facilities, and that this label is mapped to a (normalized, CSCS risk-inverted) comfort score $h(w_a^{\mathrm{b}})$. In this context, we seek node-, arc-, and path-level metrics whose meaning is clear on a physical cycling network, for instance capturing safety and comfort for intersections, roads and paths. However, a central methodological requirement is monotonicity with respect to DoS upgrades. Suppose that an optimization model chooses a set of link improvements of the form

$$w_a^{\mathrm{b}} \longrightarrow \tilde{w}_a^{\mathrm{b}} \geq w_a^{\mathrm{b}}, \qquad a \in \mathcal{A}^{\mathrm{b}},$$

which, by monotonicity of h, implies $h(\tilde{w}_a^{\mathrm{b}}) \geq h(w_a^{\mathrm{b}})$ for every arc. Let $J(\boldsymbol{w}^{\mathrm{b}})$ denote a scalar network-level indicator constructed from a given centrality metric (for example, the strength or closeness), computed on the DoS-weighted network that is a candidate for measuring safety/comfort on the cycling network. For J to behave consistently as an optimization objective, it must react to upgrades in a single direction across all arcs upgrades. That is, either

$$J(\tilde{\boldsymbol{w}}^{\mathbf{b}}) \;\geq\; J(\boldsymbol{w}^{\mathbf{b}}) \quad \text{for every feasible upgrade (non-decreasing behavior),} \tag{3.15}$$

or

$$J(\tilde{\boldsymbol{w}}^{\mathbf{b}}) \;\leq\; J(\boldsymbol{w}^{\mathbf{b}}) \quad \text{for every feasible upgrade (non-increasing behavior).} \tag{3.16}$$

If increasing the DoS of one link raises the value of J while upgrading another link lowers it, then the metric fails to preserve the physical ordering implied by the DoS scale. Such behavior makes the metric unsuitable for optimization: the algorithm could be penalized for making parts of the network safer, despite all upgrades moving in the same direction in terms of real-world comfort and safety. This monotonicity requirement has direct implications for classical centrality measures. Metrics defined purely on topology, such as unweighted degree, closeness, betweenness, or eigenvector centrality, do not respond to DoS changes, because arc labels do not affect the underlying graph structure. To make them meaningfully sensitive to DoS, we adopt the weighted versions above: unweighted degree becomes node strength $c_{\mathrm{str}}(v)$, while closeness and betweenness are computed using effective distances $\ell(a)$ derived from w_a (or equivalently from $h(w_a)$ through a decreasing cost transform). Accessibility-based metrics rely on random-walk transition weights and thus also respond to these DoS-derived costs.

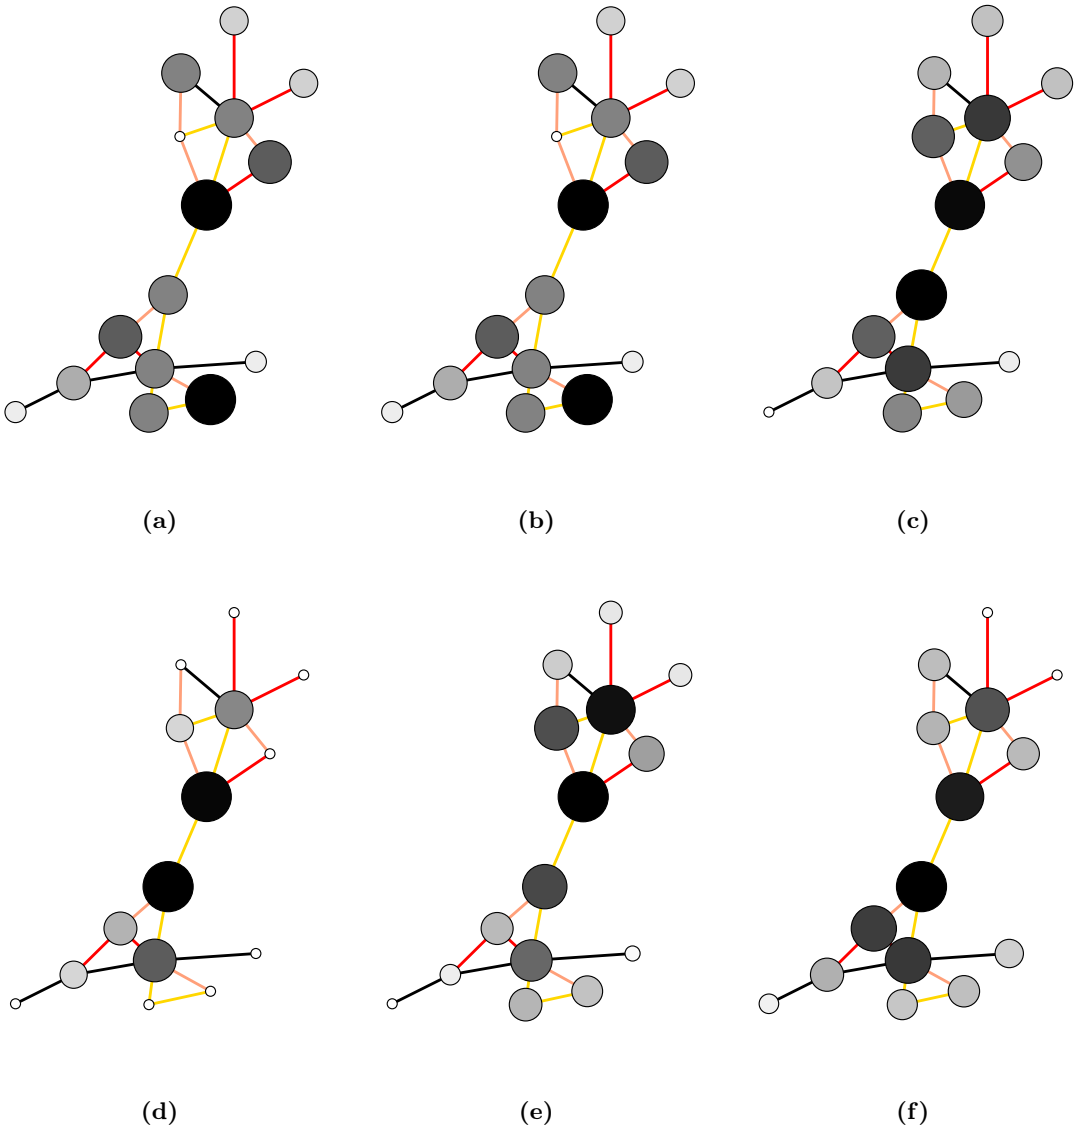

Figure 3.3: Comparison of DoS-weighted centrality measures: a, strength centrality. b, k-core. c, closeness centrality. d, betweenness centrality e, eigenvector centrality. f, accessibility. The size and darkness of each node are proportional to its centrality measure and the color of the arcs is according to its DoS: 6: ●, 5: ●, 4: ●, 3: ●, 2: ●, 1: ●.

Why insist on monotonicity in this chapter? The monotonicity requirement is not meant as a claim that non-monotone metrics are useless in general. Rather, it is a modeling choice tied to the role of the metric in this chapter: we use graph indicators as optimization objectives for upgrade decisions that, by construction, represent infrastructure improvements (higher DoS $\Rightarrow$ higher separation and higher CSCS weight). In this setting, monotonicity provides a basic consistency guarantee: if all upgrades move in the same physical direction (more separation), the network-level score used for optimization should not penalize an im-

provement purely because of how the metric aggregates information. Without this property, the optimization can become difficult to interpret, because a "better" DoS configuration may yield a lower objective value even when every modified arc is upgraded.

This does not contradict the fact that non-local metrics such as betweenness, eigenvector centrality, or accessibility capture important structural aspects of networks. The issue is interpretability as an upgrade objective: when these metrics can increase or decrease under different upgrades, it becomes unclear whether a change in the score reflects a genuine degradation of cycling comfort/safety or merely a shift in global network structure that the metric is sensitive to. In other words, non-monotonicity does not invalidate the metric as a descriptive tool, but it weakens its direct use as a proxy for "more DoS implies more perceived safety/comfort" when embedded in an optimization loop.

Finally, it is important to acknowledge that network interventions can produce counter-intuitive system effects, as illustrated by Braess-type paradoxes in traffic networks: improving a link can worsen equilibrium performance when users re-route strategically. Our monotonicity analysis does not claim that all outcomes improve after upgrades; it only ensures that the particular graph-theory proxy used in this chapter is aligned with the intended direction of infrastructure improvement under fixed network topology. Capturing the possibility of adverse network-wide effects requires a behavioral layer (route choice, congestion, and multi-modal interactions), which motivates the flow-aware and multi-modal formulations developed in subsequent chapters.

3.2.3 Numerical Monotonicity Test

To evaluate whether these metrics exhibit the required monotonicity, we conducted a numerical test in which random DoS upgrades were sampled on academic example of Fig. 3.3, and the resulting change in the network-level value of each metric was recorded. In each trial, we first sampled a DoS configuration $\boldsymbol{w}^{\mathrm{b}} = \left[w_a^{\mathrm{b}}\right]_{a\in\mathcal{A}^{\mathrm{b}}}$ by assigning to every arc a random value in $\{1,\dots,6\}$. We then selected a random arc $a^\star$ and upgraded its label to $\tilde{w}^{\mathrm{b}}_{a^\star} = \min\{w^{\mathrm{b}}_{a^\star}+1, 6\}$, keeping all other w_a^{b} unchanged, which in turn updates the associated comfort scores via $h(w_a^{\mathrm{b}}) \mapsto h(\tilde{w}_a^{\mathrm{b}})$. For each centrality metric, we computed the corresponding network-level indicator $J(\boldsymbol{w}^{\mathbf{b}})$ and $J(\tilde{\boldsymbol{w}}^{\mathbf{b}})$ (e.g., average node strength or average closeness), and recorded whether J increased ("Up"), decreased ("Down"), or remained unchanged ("Zero"). Fig. 3.4 summarizes the outcomes over 1000 such random upgrades.

The results reveal a clear pattern: only two metrics satisfy a consistent monotonic response to DoS upgrades. Strength centrality is perfectly non-decreasing, while closeness centrality is also non-decreasing, except in cases where the upgraded link does not affect any shortest-path distances (leading to a "Zero" change). All remaining metrics exhibit mixed behavior, with both increases and decreases under DoS upgrades, reflecting their non-local and often nonlinear dependence on shortest paths or walk-based transition probabilities. Since our aim is to use graph metrics both as predictors of perceived safety and comfort and as quantities with a clear, physical interpretation on cycling networks, we retain only these two centralities. In

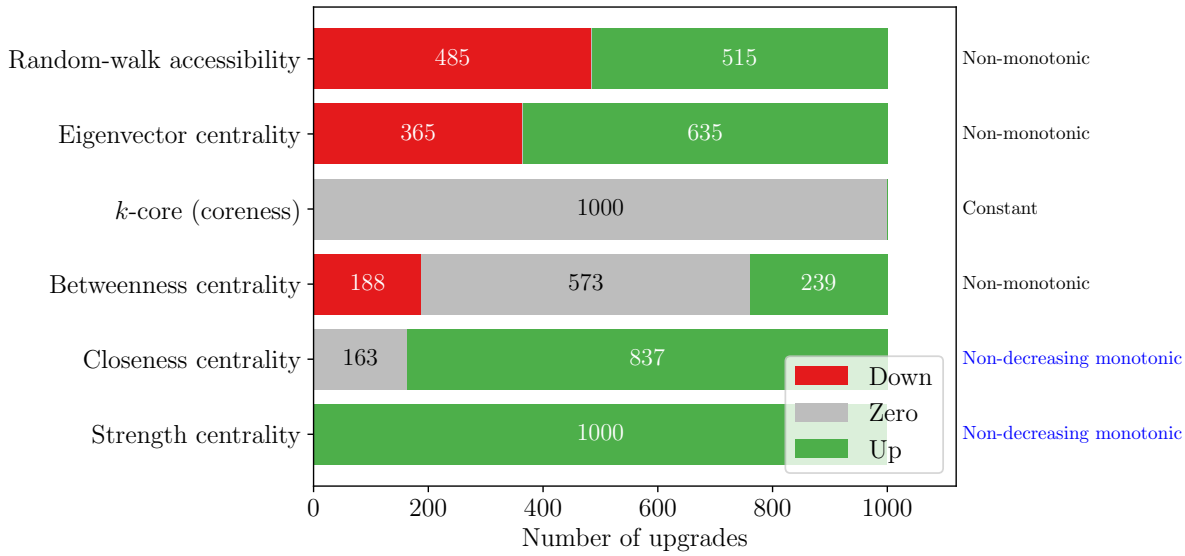

Figure 3.4: Monotonicity behavior of weighted graph metrics under 1000 random DoS upgrades. "Up" indicates that the graph metric increased after the upgrade, "Down" that it decreased, and "Zero" that it remained unchanged.

the remainder of this section, we therefore specialize strength centrality to define a node-level indicator of safety and comfort at intersections, construct an arc-based closeness centrality on the line graph to characterize the embedding of individual links within safe routes, and introduce a new path-level metric that captures the extent and continuity of high-DoS corridors. All three are designed to be monotone non-decreasing under DoS upgrades and to admit intuitive interpretations on real networks.

3.2.4 Final Metrics for Predicting Perceived Safety and Comfort

The previous analysis highlights that not all classical weighted graph metrics behave consistently under DoS upgrades. In particular, node strength (used here as a node-level metric) is guaranteed to be monotone non-decreasing, and the arc-based closeness measure introduced in this chapter is also monotone non-decreasing analogous to its node version. We therefore adopt these two indicators as the core node and link-level metrics for modeling cycling safety and comfort in a DoS-weighted network. To enable meaningful comparisons across networks and upgrade scenarios, both metrics are normalized with respect to a fully safe–comfortable (FS) reference network in which all links attain the maximum DoS level. Finally, since cyclists also value the continuity of safe infrastructure along routes, we complement these local indicators with a path-level metric that quantifies the extent of connected high-DoS corridors through the size of the largest high-DoS connected component.

Node Strength

In weighted graphs, the analogous concept to node degree is node strength, which quantifies the strength of connections a node possesses. Here, since arc labels correspond to DoS levels w_a, node strength provides a local measure where a high value indicates a safer and more comfortable intersection for cyclists. Given $v \in \mathcal{V}^{\mathrm{b}}$, analogous to Eq. (3.6) its local strength is defined as:

$$s(v) = \sum_{a \in \mathcal{N}_v} h(w_a^{\mathrm{b}}). \tag{3.17}$$

with $\mathcal{N}_v \subset \mathcal{A}^{\mathrm{b}}$ the set of adjacent arcs to v. The normalized local strength of v is defined as:

$$s_N(v) = \frac{s(v)}{s_{FS}(v)} \in [0,1]\,. \tag{3.18}$$

The normalized global strength $s_{\mathcal{G}^{\mathrm{b}}}$ is then defined as:

$$s_{\mathcal{G}}^{\mathrm{b}} = \frac{1}{|\mathcal{V}^{\mathrm{b}}|} \sum_{v \in \mathcal{V}^{\mathrm{b}}} s_N(v) \in [0,1]\,, \tag{3.19}$$

where $|\mathcal{V}^{\mathrm{b}}|$ indicates the cardinality of $\mathcal{V}^{\mathrm{b}}$.

Arc-based Closeness Centrality on the Line Graph

While $s(v)$ is a node-level metric, we also require an arc-based indicator that reflects how well each link is embedded within a system of safe, comfortable routes. To this end, we build an arc-based analogue of closeness centrality using the line graph $L(\mathcal{G}^{\mathrm{b}})$ of $\mathcal{G}^{\mathrm{b}}$. The line graph $L(\mathcal{G}^{\mathrm{b}})$ is defined as follows: its set of nodes is the original set of arcs, $V(L(\mathcal{G}^{\mathrm{b}})) = \mathcal{A}^{\mathrm{b}}$. For every pair of incident arcs $a, a' \in \mathcal{A}^{\mathrm{b}}$, we create an arc $e_{a,a'} = (e_a, e_{a'}) \in \mathcal{A}^{\mathrm{b}}(L(\mathcal{G}^{\mathrm{b}}))$. We assign to $e_{a,a'}$ an effective weight and a geographical length by

$$w_{e_{a,a'}} = \frac{1}{2}\big(h(7 - w_a^{\mathrm{b}}) + h(7 - w_{a'}^{\mathrm{b}})\big), \qquad d_{e_{a,a'}} = \frac{1}{2}(d_a^{\mathrm{b}} + d_{a'}^{\mathrm{b}}), \tag{3.20}$$

so that transitions between high-DoS arcs (large w_a^{b} and $w_{a'}^{\mathrm{b}}$) are associated with smaller penalties, because $h(7 - w_a^{\mathrm{b}})$ and $h(7 - w_{a'}^{\mathrm{b}})$ decrease as DoS increases. We then define the effective cost of $e_{a,a'}$ as

$$\tilde{\ell}_{e_{a,a'}} = d_{e_{a,a'}} \cdot w_{e_{a,a'}}. \tag{3.21}$$

For $e_a \in \mathcal{A}^{\mathrm{b}}$, denote by $\tilde{d}(e_a, e_{a'})$ the length of a shortest path from e_a to $e_{a'}$ in $L(\mathcal{G}^{\mathrm{b}})$ with respect to costs $\tilde{\ell}_{e_{a,a'}}$. The local (arc-based) closeness of e_a is then defined as

$$c(e_a) = \frac{|\mathcal{A}^{\mathrm{b}}| - 1}{\sum_{e_{a''} \in \mathcal{A}^{\mathrm{b}}} \tilde{d}(e_a, e_{a''})}, \tag{3.22}$$

with $\tilde{d}(e_a, e_a) = 0$ for all $e_a \in \mathcal{A}^{\mathrm{b}}$. If $L(\mathcal{G}^{\mathrm{b}})$ is disconnected and e_a and $e_{a'}$ belong to different connected components, we set $\tilde{d}(e_a, e_{a'}) = 0$ and if unreachable, we define $c(e_a) = 0$. As

before, we introduce a fully safe–comfortable reference network and define a normalized local arc closeness

$$c_N(e_a) = \frac{c(e_a)}{c_{FS}(e_a)} \in [0,1], \tag{3.23}$$

and a global indicator

$$c_{\mathcal{G}^{\mathrm{b}}} = \frac{1}{|\mathcal{A}^{\mathrm{b}}|} \sum_{e_a \in \mathcal{A}^{\mathrm{b}}} c_N(e_a) \in [0,1]. \tag{3.24}$$

Because DoS upgrades decrease the effective costs along paths in $L(\mathcal{G}^{\mathrm{b}})$, the arc-based closeness $c(e_a)$ and its normalized version $c_{\mathcal{G}^{\mathrm{b}}}$ are empirically observed to be monotone non-decreasing in our experiment, making them suitable arc-level metrics.

Largest Connected Component of High-DoS Corridors

Finally, we introduce a path-level metric based on the extent of connected corridors with high DoS. The underlying assumption is that cyclists value networks where high-DoS facilities are organized into long connected paths, rather than being scattered as isolated segments among low-DoS roads. For a given DoS threshold $\gamma \in \{2, \ldots, 6\}$, let

$$C_\gamma = \{a \in \mathcal{A}^{\mathrm{b}} : w_a^{\mathrm{b}} \geq \gamma\} \tag{3.25}$$

denote the set of arcs in $\mathcal{G}^{\mathrm{b}}$ with DoS at least γ, and let $\mathcal{H}_\gamma$ be the largest connected subgraph of $\mathcal{G}^{\mathrm{b}}$ induced by C_γ. We define a normalized length of this high-DoS subgraph as

$$l_N(\gamma) = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(\mathcal{H}_\gamma)} d_a^{\mathrm{b}} \cdot h(w_a^{\mathrm{b}})}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(\mathcal{H}_{FS})} d_a^{\mathrm{b}} \cdot h(w_a^{\mathrm{b}})} \in [0,1], \tag{3.26}$$

where $\mathcal{H}_{FS}$ is the fully safe–comfortable (FS) network in which all arcs have $w_a = 6$. The denominator is therefore a positive constant that depends only on $\mathcal{G}^{\mathrm{b}}$ and not on w or γ. We then aggregate across thresholds using exponentially decaying weights that emphasize higher DoS levels:

$$l_{\mathcal{G}^{\mathrm{b}}} = \sum_{\gamma=2}^{6} \frac{\exp(\gamma)}{\alpha_1} l_N(\gamma) \in [0,1], \tag{3.27}$$

where

$$\alpha_1 = \sum_{i_1=2}^{6} \exp(i_1) \tag{3.28}$$

ensures proper normalization.

The three indicators $s_{N,\mathcal{G}^{\mathrm{b}}}$, $c_{\mathcal{G}^{\mathrm{b}}}$, and $l_{\mathcal{G}^{\mathrm{b}}}$ thus provide complementary perspectives on network quality: a node-based measure of local protection, an arc-based measure of embeddedness within safe routes, and a path-based measure of continuous high-DoS corridors. All three are constructed to be monotone non-decreasing under DoS upgrades, and we hypothesise that

they are positively correlated with perceived safety and comfort. This makes them natural candidates as explanatory features in our safety–comfort prediction model.

Theorem 3.1: (Monotonicity of the proposed network metrics).

Let $\mathcal{G}^{\mathrm{b}} = (\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}})$ be a fixed directed graph, and let $\boldsymbol{w}^{\mathrm{b}} = [w_a^{\mathrm{b}}]_{a \in \mathcal{A}^{\mathrm{b}}}$ and $\tilde{\boldsymbol{w}}^{\mathrm{b}} = [\tilde{w}_a^{\mathrm{b}}]_{a \in \mathcal{A}^{\mathrm{b}}}$ be two DoS assignments such that

$$\tilde{w}_a^{\mathrm{b}} \geq w_a^{\mathrm{b}} \quad \text{for all } a \in \mathcal{A}^{\mathrm{b}},$$

with $w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}} \in \{1, \ldots, 6\}$. Let $h : \{1, \ldots, 6\} \to \mathbb{R}_{>0}$ be a strictly increasing function mapping each DoS level to its CSCS values. Assume that every arc length $d_a^{\mathrm{b}} > 0$, and that effective arc costs are constructed as $\ell(a) = d_a^{\mathrm{b}} \cdot f(w_a^{\mathrm{b}})$ for some strictly decreasing function $f \colon \{1, \ldots, 6\} \to \mathbb{R}_{>0}$. Let $s_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$, $c_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$, and $l_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$ denote, respectively, the normalized global strength, the global arc-based closeness on the line graph, and the largest connected component metric defined in Section 3.2.4–3.2.4. Then:
i) $s_{\mathcal{G}^{\mathrm{b}}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq s_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$ (with strict inequality if at least one $\tilde{w}_a > w_a$);
ii) $c_{\mathcal{G}^{\mathrm{b}}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq c_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$;
iii) $l_{\mathcal{G}^{\mathrm{b}}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq l_{\mathcal{G}^{\mathrm{b}}}(\boldsymbol{w}^{\mathrm{b}})$.

In particular, all three metrics are monotone non-decreasing with respect to DoS upgrades.

Proof. The proof of the Theorem is shown in Appendix A. □

3.3 Model Design

Having identified the graph-theory metrics that best reflect perceived safety and comfort, and having verified that the selected indicators satisfy the required *monotonicity* property under DoS upgrades, we now proceed to the model design. In this section, we formalize the decision variables and upgrade actions on the cycling network, define the objective functions based on the retained metrics, and introduce the budget constraints that govern feasible interventions. These elements are then assembled into the optimization framework used throughout the remainder of this chapter to compute upgrade plans that improve safety and comfort in a principled and comparable way.

Let $\tilde{M} \in \mathbb{R}^{n \times 3}$ be the matrix whose columns are the metrics $s_{\mathcal{G}^{\mathrm{b}}}, c_{\mathcal{G}^{\mathrm{b}}}, l_{\mathcal{G}^{\mathrm{b}}}$, and where n is the total number of cities (samples). The vectors $y_s \in \mathbb{R}^n$ and $y_c \in \mathbb{R}^n$ collect the perceived safety and comfort scores from the "Baromètre Vélo" survey. The central question in this section is to find a function $f(\tilde{M}, \theta)$ and a parameter set θ such that

$$f(\tilde{M}, \theta) \approx y, \tag{3.29}$$

where

$$f(\tilde{M},\theta)=\begin{bmatrix}f_s(\tilde{M},\theta_s)^\mathsf{T}\\ f_c(\tilde{M},\theta_c)^\mathsf{T}\end{bmatrix},\qquad y=\begin{bmatrix}y_s^\mathsf{T}\\ y_c^\mathsf{T}\end{bmatrix},\qquad \theta=\begin{bmatrix}\theta_s^\mathsf{T}\\ \theta_c^\mathsf{T}\end{bmatrix}.$$

Here, $f_s, f_c \in \mathbb{R}^n$ represent the individual predictions for perceived safety and comfort, respectively, and θ_s and θ_c are the corresponding parameter vectors.

As detailed in Section 3.2.4, it is essential that the metrics exhibit non-decreasing monotonic behavior with respect to upgrades of the cycling infrastructure. To formalize this, we introduce the set $\mathbb{H}$ of graphs that are isomorphic to $\mathcal{G}^\mathrm{b}$ (same nodes and arcs) but with DoS weights at least as large as those of $\mathcal{G}^\mathrm{b}$:

$$\mathbb{H}=\left\{\tilde{\mathcal{G}}^\mathrm{b}:\mathcal{V}^\mathrm{b}(\tilde{\mathcal{G}}^\mathrm{b})=\mathcal{V}^\mathrm{b}(\mathcal{G}^\mathrm{b}),\,\mathcal{A}^\mathrm{b}(\tilde{\mathcal{G}}^\mathrm{b})=\mathcal{A}^\mathrm{b}(\mathcal{G}^\mathrm{b}),\,\forall a\in\mathcal{A}^\mathrm{b}(\mathcal{G}^\mathrm{b}),\,w_a^\mathrm{b}\le w_a^{\tilde{\mathcal{G}}^\mathrm{b}}\right\},$$

where w_a and $w_a^{\tilde{\mathcal{G}}^\mathrm{b}}$ denote the DoS weights of arc a in $\mathcal{G}^\mathrm{b}$ and $\tilde{\mathcal{G}}^\mathrm{b}$, respectively. Then, we denote the metrics of a graph $\mathcal{G}^\mathrm{b}$ by the vector

$$m_{\mathcal{G}^\mathrm{b}}=\left[s_{\mathcal{G}^\mathrm{b}},\,c_{\mathcal{G}^\mathrm{b}},\,l_{\mathcal{G}^\mathrm{b}}\right]\in\mathbb{R}^{1\times 3}.$$

The model used to predict perceived safety and comfort can therefore be stated as follows.

Problem 3.1 (Monotonicity-constrained model fitting).

For all $\tilde{\mathcal{G}}^\mathrm{b}\in\mathbb{H}$, find g^* and θ^* such that

$$(f^*,\theta^*)=\underset{f,\theta}{\operatorname{argmin}}\left\|f(\tilde{M},\theta)-y\right\|^2 \tag{3.30}$$

$$\text{subject to}\quad f_s(m_{\mathcal{G}^\mathrm{b}},\theta_s)\ \le\ f_s(m_{\tilde{\mathcal{G}}^\mathrm{b}},\theta_s), \tag{3.31}$$

$$f_c(m_{\mathcal{G}^\mathrm{b}},\theta_c)\ \le\ f_c(m_{\tilde{\mathcal{G}}^\mathrm{b}},\theta_c), \tag{3.32}$$

where the constraints Eq. (3.31)–(3.32) enforce monotone non-decreasing behavior of the models with respect to DoS upgrades.

Problem 3.1 is a functional optimization problem. Optimizing over a finite-dimensional function class (for example, polynomials up to a given degree) is typically more tractable than working in an infinite-dimensional space of functions. We therefore restrict our attention to first-degree polynomial models and compare their performance with that of a data-driven machine learning approach.

3.3.1 Linear Regression Approach

In this section, we start with a linear regression model due to its simplicity. We then employ statistical methods such as the P-value and t-test [104] to identify the most significant metrics for predicting perceived safety and comfort. This allows us to address the linear version of Problem 3.1 using only the significant metrics. Hence, the proposed form for $f(\tilde{M},\theta)$ is

$$f(\tilde{M},\theta)=\begin{bmatrix}[\mathbf{1}\mid\tilde{M}] & 0\\ 0 & [\mathbf{1}\mid\tilde{M}]\end{bmatrix}\theta=M\theta=f(M,\theta), \tag{3.33}$$

No	City	$s_{\mathcal{G}^b}$	$c_{\mathcal{G}^b}$	$l_{\mathcal{G}^b}$	y_s	y_c
1	Angers	0.4420	0.3575	0.1113	0.420	0.484
2	Annecy	0.3469	0.2887	0.0543	0.448	0.552
3	Bordeaux	0.4233	0.3880	0.0653	0.440	0.458
4	Chambéry	0.3939	0.3291	0.1729	0.514	0.564
5	Dunkerque	0.4565	0.5049	0.0870	0.488	0.524
6	Grenoble	0.4170	0.4263	0.2028	0.608	0.674
7	La Rochelle	0.3557	0.4275	0.1529	0.566	0.608
8	Le Havre	0.5504	0.4963	0.1209	0.458	0.444
9	Lyon	0.3453	0.4078	0.1665	0.440	0.508
10	Nantes	0.3507	0.3445	0.0557	0.470	0.522
11	Paris	0.3486	0.5400	0.0537	0.412	0.462
12	Rennes	0.4433	0.3611	0.0426	0.480	0.534
13	Strasbourg	0.5326	0.6087	0.3760	0.622	0.654
14	Tours	0.3474	0.3210	0.0731	0.430	0.468

Table 3.2: Values of $s_{\mathcal{G}^b}, c_{\mathcal{G}^b}$ and $l_{\mathcal{G}^b}$ for fourteen French municipalities. Safety (y_s) and comfort (y_c) values are scaled to the interval $[0, 1]$.

where $[\mathbf{1} \mid \tilde{M}] \in \mathbb{R}^{n\times 4}$ is the original matrix $\tilde{M}$ with an additional column of ones placed at the beginning to introduce a bias term in the regression, and thus $\theta \in \mathbb{R}^8$. Solving Eq. (3.33) is equivalent to finding

$$\theta^* = \underset{\hat{\theta}\in\mathbb{R}^8}{\operatorname{argmin}} \left\| M\hat{\theta} - y \right\|^2. \tag{3.34}$$

Retain Relevant Metrics

To solve Eq. (3.34), we apply the least-squares method [105], using $\tilde{M}$ and y from Table 3.2. This enables us to compute P-values and t-statistics for each coefficient and to assess the significance of each metric in predicting safety and comfort, testing against the null hypothesis H_0 that a metric is not a significant predictor.

With a significance level of 0.05, the critical t-value is 2.22 for $df = 10$ [106]. The hypotheses are:

- H_0 : P-value > 0.05 and $|t|$-value < 2.22 (metric is insignificant);
- H_A : P-value < 0.05 and $|t|$-value > 2.22 (metric is significant).

Applying least squares to the safety and comfort models (see Table 3.3), we observe that $s_{\mathcal{G}^b}$ and $c_{\mathcal{G}^b}$ have P-values above 0.05 and $|t|$-values below 2.22, indicating that they are not statistically significant in this linear setting. Their local nature may limit their explanatory power at the city scale. By contrast, $l_{\mathcal{G}^b}$ satisfies the significance criteria for both safety and comfort, confirming its relevance as a predictor.

θ^*	Coefficient	$\lvert t\rvert$-value	P-value
Safety			
Intercept	0.4255	4.816	0.001
$s_{\mathcal{G}^{\mathrm{b}}}$	-0.0063	0.026	0.980
$c_{\mathcal{G}^{\mathrm{b}}}$	-0.0330	0.173	0.866
$l_{\mathcal{G}^{\mathrm{b}}}$	0.6149	3.340	0.007
Comfort			
Intercept	0.5940	5.980	0.000
$s_{\mathcal{G}^{\mathrm{b}}}$	-0.2311	0.848	0.416
$c_{\mathcal{G}^{\mathrm{b}}}$	-0.1316	0.614	0.553
$l_{\mathcal{G}^{\mathrm{b}}}$	0.7112	3.437	0.006

Table 3.3: Regression results for solving Eq. (3.34) for safety and comfort.

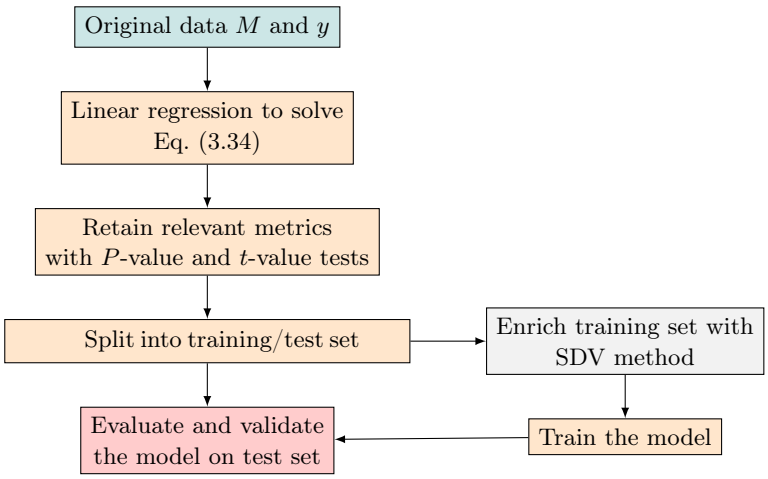

Figure 3.5: Flowchart of the model training and testing process.

Synthetic Data Vault

In current practice, data are typically divided into training and validation sets using a ratio of 75% to 25%. After removing the columns $s_{\mathcal{G}^{\mathrm{b}}}$ and $c_{\mathcal{G}^{\mathrm{b}}}$ and selecting 10 random samples for training for both safety and comfort, the matrix $M \in \mathbb{R}^{20\times 4}$, the vector $\theta \in \mathbb{R}^4_{\geq 0}$ in Eq. (3.33), and the target vector $y \in \mathbb{R}^{20}$ in Eq. (3.29) are redefined according to the 10 selected cities (stacking safety and comfort). The remaining 4 cities are used for validation.

To improve the conditioning of M under the small-sample setting (10 training cities, stacked for safety and comfort), we use the Synthetic Data Vault (SDV) generator [107] to create additional synthetic observations. Concretely, we fit an SDV model on the training set and sample $N_{\mathrm{SDV}} = 1000$ additional rows, yielding an augmented design matrix M_{aug} and target vector y_{aug}. The purpose is not to add new information, but to stabilize the regres-

sion numerically by densifying the empirical distribution of the observed points and reducing variance in the estimated coefficients. We verified that results are stable for N_{SDV} in the range $[500, 2000]$ and we report the best-performing setting on the validation cities. We use $N_{\text{SDV}} = 1000$ synthetic observations because the original training design matrix is small and can be poorly conditioned; augmenting the sample stabilizes coefficient estimation and reduces sensitivity to the particular train/validation split.

Lemma 3.1: (When NNLS coincides with unconstrained least squares).

Let $\hat{\theta}^\star$ be a minimizer of the unconstrained least-squares problem

$$\hat{\theta}^\star \in \arg\min_{\hat{\theta}\in\mathbb{R}^4} \|M\hat{\theta} - y\|^2.$$

If $\hat{\theta}^\star \in \mathbb{R}^4_{\geq 0}$, then $\hat{\theta}^\star$ is also a minimizer of the non-negative least-squares problem

$$\hat{\theta}^\star \in \arg\min_{\hat{\theta}\in\mathbb{R}^4_{\geq 0}} \|M\hat{\theta} - y\|^2.$$

Proof. If $\hat{\theta}^\star \in \mathbb{R}^4_{\geq 0}$, then it is feasible for the NNLS problem. Since the NNLS feasible set $\mathbb{R}^4_{\geq 0}$ is a subset of $\mathbb{R}^4$, no feasible point can attain a strictly smaller least-squares value than the global minimum over $\mathbb{R}^4$. Therefore, $\hat{\theta}^\star$ is also optimal for the NNLS problem. □

3.3.2 Model complexity and evaluation criteria.

Because the number of observations is limited, model comparison should not rely only on in-sample fit. In particular, models with additional predictors may appear to perform better simply because they are more flexible, even if they do not generalize better. To address this issue, we focus on two complementary criteria: the corrected Akaike Information Criterion (AICc) and the leave-one-out cross-validated root mean squared error (LOOCV-RMSE).

The AICc is defined as

$$\text{AICc} = \text{AIC} + \frac{2p(p+1)}{n-p-1},$$

where n is the number of observations and p is the number of estimated parameters, including the intercept. AICc is particularly appropriate in small-sample settings because it penalizes overly complex models more strongly than the standard AIC. Lower values indicate a better compromise between goodness of fit and parsimony. The LOOCV-RMSE is defined as

$$\text{LOOCV-RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i^{(-i)}\right)^2},$$

where $\hat{y}_i^{(-i)}$ denotes the prediction for observation i obtained from a model fitted without that observation. This criterion measures predictive performance on unseen data and therefore

Table 3.4: Comparison of the most relevant models using AICc and LOOCV-RMSE. Lower values are better for both criteria.

Target	Predictors	AICc	LOOCV-RMSE
y_s	$l_{\mathcal{G}^\mathrm{b}}$	-44.4733	0.0471
y_s	$c_{\mathcal{G}^\mathrm{b}}$, $l_{\mathcal{G}^\mathrm{b}}$	-41.2221	0.0490
y_s	$s_{\mathcal{G}^\mathrm{b}}$, $c_{\mathcal{G}^\mathrm{b}}$, $l_{\mathcal{G}^\mathrm{b}}$	-37.1786	0.0562
y_c	$l_{\mathcal{G}^\mathrm{b}}$	-38.8475	0.0582
y_c	$c_{\mathcal{G}^\mathrm{b}}$, $l_{\mathcal{G}^\mathrm{b}}$	-36.9737	0.0562
y_c	$s_{\mathcal{G}^\mathrm{b}}$, $l_{\mathcal{G}^\mathrm{b}}$	-37.4266	0.0598

provides an indication of the model's ability to generalize. Lower values indicate better predictive accuracy. The two criteria are complementary: AICc evaluates whether a gain in fit justifies the increase in model complexity, while LOOCV-RMSE evaluates whether the model actually predicts well outside the fitted sample.

For safety, the model using only $l_{\mathcal{G}^\mathrm{b}}$ is preferred according to both criteria. It has the lowest AICc (-44.4733) and the lowest LOOCV-RMSE (0.0471), which indicates that it provides the best balance between parsimony and predictive performance. Adding further predictors does not improve the model sufficiently to justify the increase in complexity. For comfort, the model based only on $l_{\mathcal{G}^\mathrm{b}}$ has the lowest AICc (-38.8475), which makes it the preferred model from the viewpoint of complexity-adjusted fit. The model $(c_{\mathcal{G}^\mathrm{b}}, l_{\mathcal{G}^\mathrm{b}})$ achieves a slightly lower LOOCV-RMSE (0.0562 versus 0.0582), suggesting a small gain in predictive accuracy. However, this improvement is limited and comes at the cost of a more complex specification. Overall, the results still favor the parsimonious model based solely on $l_{\mathcal{G}^\mathrm{b}}$.

In summary, AICc and LOOCV-RMSE both indicate that $l_{\mathcal{G}^\mathrm{b}}$ is the most informative graph-theoretic predictor for safety and comfort, while the additional metrics provide only limited extra benefit.

3.3.3 Linear Regression vs. Ensemble ML Models

We fit a simple linear predictor using non-negative least squares (NNLS) with $l_{\mathcal{G}^\mathrm{b}}$ as the sole explanatory variable. This choice is motivated by the significance analysis in Table 3.3, where $l_{\mathcal{G}^\mathrm{b}}$ is the only metric that consistently explains perceived safety and comfort. The estimated coefficient vector (stacking safety and comfort) is

$$\tilde{\theta} = [0.41,\ 0.58,\ 0.45,\ 0.57]^\top . \tag{3.35}$$

Model complexity and evaluation in a small-sample regime. Table 3.2 contains only $n = 14$ city observations, so model complexity must be handled carefully to avoid overfitting. In our linear specification, each target (safety and comfort) is modeled as an affine function of $l_{\mathcal{G}^\mathrm{b}}$, i.e., one intercept and one slope per target. This corresponds to $p = 2$ fitted parameters per output (and $p = 4$ parameters in total when the two outputs are stacked). In addition,

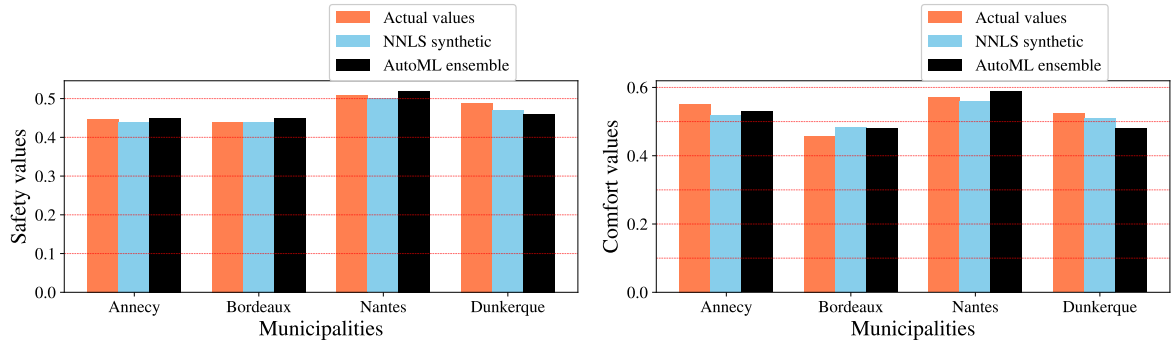

Figure 3.6: Predicted values versus the actual values of perceived safety (right) and comfort (left). The average error is 2.23 % and 3.15 % for Auto-ML Ensemble and NNLS, respectively, for safety, and 5.99 % and 7.36 % for Auto-ML Ensemble and NNLS, respectively, for comfort.

unconstrained least squares can yield negative coefficients in such a small dataset; we therefore use NNLS to enforce $\hat{\theta} \geq 0$, preserving the intended interpretation that increasing $l_{\mathcal{G}^b}$ should not decrease predicted safety/comfort under DoS upgrades.

Validation against an AutoML ensemble. To validate this simple model against a flexible predictor, we also train an AutoGluon-Tabular ensemble [108]. AutoGluon fits a library of base models (e.g., boosted trees, random forests, k-NN, and a neural network) and combines them through bagging and multi-layer stacking; the final predictor is obtained by ensemble selection, which learns a weighted combination of candidate models to reduce variance and mitigate overfitting. Fig. 3.6 compares predicted and observed values. Using the average absolute error (in percent) as the evaluation criterion, the AutoML ensemble achieves 2.23% error for safety and 5.99% for comfort, while NNLS achieves 3.15% for safety and 7.36% for comfort. Given the small difference in accuracy and the need for an interpretable model that can be embedded in the optimization formulations, we retain NNLS as the main predictive layer in the remainder of this thesis.

Network Optimal Design

In this section, we formulate the network-upgrade optimization problem using the identified linear regression model. Upgrading a road segment corresponds to increasing its DoS level, which in turn increases its CSCS value $h(w_a^b)$ and incurs a cost per meter. DoS levels can only increase within the discrete set $\{1, \ldots, 6\}$. The goal is to select arcs for upgrade so as to maximize the expected improvement in perceived safety and comfort in a given cycling network $\mathcal{G}^b$, while respecting a fixed budget B.

Let $\boldsymbol{w}^{\mathbf{b}} = [w_a^b]_{a \in \mathcal{A}^b}$ denote the original DoS assignment on $\mathcal{G}^b$, and let $\tilde{\boldsymbol{w}}^b = [\tilde{w}_a^b]_{a \in \mathcal{A}^b}$ be a candidate upgraded configuration. Feasible upgrades are collected in the set

$$\Omega = \left\{ \tilde{\boldsymbol{w}}^b \in \{1, \ldots, 6\}^{|\mathcal{A}^b|} : w_a^b \leq \tilde{w}_a^b \leq 6, \ \forall a \in \mathcal{A}^b \right\},$$

i.e. every arc can only be upgraded to an equal or higher DoS level, up to the maximum DoS 6. The corresponding CSCS weights are then given by $h(w_a^{\mathrm{b}})$ and $h(\tilde{w}_a^{\mathrm{b}})$ for the original and upgraded networks, respectively. Based on the linear model structure derived above and the identified parameter vector $\tilde{\theta}$ (see Eq. (3.35)), we define two predicted outputs, corresponding to the original and upgraded networks, respectively:

$$\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} = M_{\boldsymbol{w}^{\mathrm{b}}}\tilde{\theta}, \qquad \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} = M_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\tilde{\theta},$$

where $M_{\boldsymbol{w}^{\mathrm{b}}} \in \mathbb{R}^{2\times 4}$ and $M_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} \in \mathbb{R}^{2\times 4}$ have the same block structure as M in Eq. (3.33), but with only the intercept and the path-level metric $l_{\mathcal{G}^{\mathrm{b}}}$ retained (as suggested by the P-value analysis in the previous section). They are defined as

$$M_{\boldsymbol{w}^{\mathrm{b}}} = \begin{bmatrix} m_{\boldsymbol{w}^{\mathrm{b}}} & 0 \\ 0 & m_{\boldsymbol{w}^{\mathrm{b}}} \end{bmatrix}, \qquad M_{\tilde{w}^{\mathrm{b}}} = \begin{bmatrix} m_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} & 0 \\ 0 & m_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} \end{bmatrix}, \tag{3.36}$$

with row vectors $m_{\boldsymbol{w}^{\mathrm{b}}}, m_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} \in \mathbb{R}^{1\times 2}$ given by

$$m_{\boldsymbol{w}^{\mathrm{b}}} = \begin{bmatrix} 1, \ l_{\mathcal{G}^{\mathrm{b}}} \end{bmatrix}, \qquad m_{\tilde{\boldsymbol{w}}} = \begin{bmatrix} 1, \ l_{\mathcal{G}^{\mathrm{b}}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}} \end{bmatrix}, \tag{3.37}$$

where $l_{\mathcal{G}^{\mathrm{b}}_{\tilde{\boldsymbol{w}}}}$ denotes the path-level metric computed on $\mathcal{G}^{\mathrm{b}}$ with the DoS configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$ (and corresponding CSCS weights $h(\tilde{w}_a^{\mathrm{b}})$). Let $\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$ denote the cost per meter of upgrading arc a from DoS w_a^{b} to $\tilde{w}_a^{\mathrm{b}}$, and let d_a^{b} be its geometric length. The network-upgrade optimization problem can then be written as follows.

Problem 3.2 (Network-upgrade optimization problem).

$$\tilde{\boldsymbol{w}}^{\mathrm{b}^*} = \underset{\tilde{\boldsymbol{w}}^{\mathrm{b}}\in\Omega}{\operatorname{argmax}} \ \|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\| \tag{3.38}$$

$$= \underset{\tilde{\boldsymbol{w}}^{\mathrm{b}}\in\Omega}{\operatorname{argmax}} \ \left\| \left(M_{\boldsymbol{w}^{\mathrm{b}}} - M_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\right)\tilde{\theta} \right\| \tag{3.39}$$

$$\text{subject to} \quad \sum_{a\in\mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \ \leq \ B. \tag{3.40}$$

Problem 3.2 has objective and budget constraint are nonlinear and exhibit a knapsack-like combinatorial structure. In preliminary experiments on real cycling networks (Section 3.5), standard off-the-shelf solvers were either prohibitively slow or failed to converge on large instances, so we proceed to create an ad hoc algorithm due to the complexity of the task.

Lemma 3.2: (Network-upgrade optimization problem complexity).

Problem 3.2 is NP-hard with an exhaustive-search complexity of $\mathcal{O}(6^{|\mathcal{A}^{\mathrm{b}}|})$.

Proof. The proof of the Theorem is shown in Appendix A. □

To obtain tractable solutions, we therefore introduce a dedicated heuristic, the MLCC (Maximize Largest Connected Component) algorithm. MLCC is designed to approximate Problem 3.2 by prioritizing upgrades that expand the largest connected subnetwork of high-DoS facilities, while respecting the budget. The following lemma states the worst-case time complexity of the MLCC procedure. For simplicity, we denote its output upgrade plan by $\tilde{\boldsymbol{w}}^{\mathrm{b}^*}$.

Lemma 3.3: (MLCC complexity).

The MLCC algorithm runs in $\mathcal{O}\big(|\mathcal{A}^{\mathrm{b}}|\,|\mathcal{V}^{\mathrm{b}}|^4\big)$ time.

Proof. The proof of the Theorem is shown in Appendix A. □

3.4 Solution Algorithms

To solve the proposed optimization problem, we investigate four algorithmic approaches. We first introduce an ad hoc heuristic based on maximizing the largest connected component of high-quality infrastructure (MLCC). This choice is motivated by both modeling and planning considerations. Although Problem 3.2 has a knapsack-like budget constraint, its objective is not separable across arcs: upgrading a single link can change network-level structure (connectivity and corridor continuity), so the marginal benefit of an upgrade depends on which other upgrades are selected. In such settings, simple greedy rules based on a fixed benefit–cost ratio can be misleading, and dynamic-programming approaches become impractical at city scale because they would require discretizing the budget and exploring a large combinatorial state space. More importantly, MLCC reflects a widely recognized deficiency of real cycling networks: high DoS infrastructure is often fragmented into disconnected "islands" [109], and a key planning task is to remove barriers and create continuous corridors that allow end-to-end trips without forcing riders onto high-stress links. In that sense, MLCC provides an interpretable baseline aligned with practice, targeting continuity of the safest facilities rather than isolated local improvements. We then consider a standard genetic algorithm (GA) as a general-purpose metaheuristic for searching the upgrade space. Next, we propose a warm-start variant, denoted GA-U, whose initial population is seeded using the solution returned by the MLCC method. Finally, we evaluate a Bayesian optimization approach as a sample-efficient alternative to exploring the decision space under a limited evaluation.

MLCC algorithm: The idea of the MLCC (Maximize Largest Connected Component) algorithm is to systematically upgrade arcs so as to expand sub-networks that form the largest high-quality corridors. In line with the regression analysis in Section 3.3.3, the algorithm focuses on increasing the path-level metric $l_{\mathcal{G}^{\mathrm{b}}}$, which was identified as the only metric with a significant impact on predicted perceived safety and comfort. We work with DoS labels $w_a^{\mathrm{b}} \in \{1, \ldots, 6\}$ on arcs $a \in \mathcal{A}^{\mathrm{b}}$. To define shortest paths that prioritize safer infrastructure,

we introduce a DoS-based shortest path: this is the shortest path with respect to an effective length of the form

$$\ell(a) = d_a^{\mathrm{b}} \cdot g(w_a^{\mathrm{b}}), \tag{3.41}$$

where d_a^{b} is the geometric length of arc a and $g(\cdot)$ is a decreasing function of the DoS level (in practice, proportional to the complement $7 - w_a^{\mathrm{b}}$, so that higher DoS leads to a smaller effective cost). At each step of the algorithm, upgrades are performed subject to the budget constraint; once the budget is exhausted, the process stops. The output is the list of upgraded DoS levels w_a^{b} for all arcs $a \in \mathcal{A}^{\mathrm{b}}$.

For each connected component of the underlying graph, the algorithm proceeds as follows:

1) *Build a single largest DoS-6 path.* Identify the largest connected component whose arcs have the highest DoS level currently present in the network.

   1.1) If this level is strictly less than DoS 6, upgrade all arcs in this component to DoS 6.

   1.2) Otherwise, if there are several connected components already at DoS 6, connect them starting from the largest to the smallest by upgrading the arcs along the DoS-based shortest paths between components to DoS 6.

   After this step, we obtain a single connected sub-network of DoS-6 arcs, denoted by $\widehat{C}$.

2) *Enlarge $\widehat{C}$ with lower-DoS components.* Let $C_{\tilde{\gamma}}$ be the largest connected component whose arcs all have DoS level $\tilde{\gamma} < 6$ and such that $\tilde{\gamma}$ is maximal among the remaining components. Compute a DoS-based shortest path between $\widehat{C}$ and $C_{\tilde{\gamma}}$ and upgrade all arcs along this path to DoS 6. Then upgrade all arcs in $C_{\tilde{\gamma}}$ itself to DoS 6. This enlarges $\widehat{C}$ by absorbing $C_{\tilde{\gamma}}$. Repeat this step, each time selecting the next largest connected component with the highest remaining DoS level $\tilde{\gamma} < 6$, until the budget is fully utilized or all arcs reach DoS 6.

A detailed flowchart of the MLCC algorithm is provided in Appendix B.

Genetic algorithm (GA): The genetic algorithm searches of the vector of upgraded DoS levels on all arcs $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega$, where Ω is the feasible set that enforces the budget constraint (i.e., $\mathrm{cost}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \leq B$) and any DoS bounds (e.g., $w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6$). The fitness function is the objective value associated with a candidate plan:

$$F(\tilde{\boldsymbol{w}}) = \left\| \hat{y}_{\boldsymbol{w}_0^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} \right\|, \tag{3.42}$$

where $\hat{y}_{\boldsymbol{w}^{\mathrm{b}}}$ is the reference quantity under the current DoS vector $\boldsymbol{w}^{\mathrm{b}}$ and $\hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}$ is the corresponding quantity after applying upgrades $\tilde{\boldsymbol{w}}^{\mathrm{b}}$. Starting from a random feasible population, GA repeats selection, crossover, and mutation to generate new candidate plans, while retaining the best individual found so far.

Algorithm 1 Genetic algorithm GA

Input: initial DoS vector $\boldsymbol{w}^{\mathrm{b}}$; budget B
Parameters: population size n_p; generations n; mutation rate n_m; crossover rate n_c
Feasible set: $\Omega = \{\tilde{\boldsymbol{w}}^{\mathrm{b}} : \mathrm{cost}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \leq B,\ w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6\ \forall a\}$
Fitness: $F(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\|$
Output: best upgrade plan $\tilde{\boldsymbol{w}}^{\mathrm{b}}$
Initialize population $\Omega_0 \subset \Omega$ with n_p random feasible upgrade plans
Evaluate $F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ for all $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega_0$
$\tilde{\boldsymbol{w}}^{\mathrm{b}^\star} \leftarrow \arg\max_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega_0} F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$
for $t = 1$ to n **do**
    Select parent plans from Ω_0 with probability increasing in F (e.g., tournament selection)
    Create offspring by crossover (rate n_c) and mutation (rate n_m)
    Enforce feasibility: repair or discard offspring with $\mathrm{cost}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) > B$
    Update $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}$ if a better plan is found in Ω_0
end for
return $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}$

Warm-Start Genetic Algorithm (GA-U): The genetic algorithm (GA) searches for an upgrade plan $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}$ under a fixed budget B. Each individual encodes a candidate vector of upgraded DoS levels on the bicycle network and is evaluated through the corresponding objective function. The variant (GA-U) differ only in its initialization and improvement process under stagnation: GA-U targets maximize perceived safety and comfort; in addition, GA-U is warm-started by seeding its initial population with the solution produced by the MLCC heuristic (together with perturbed variants). At each generation, the population is evolved via selection, crossover, and mutation. Individuals whose total upgrade cost exceeds the available budget are deemed infeasible and discarded (or repaired) before evaluation. If no improvement is observed for a prescribed number of generations, we apply a diversification step inspired by MLCC: the current best solution is expanded by upgrading arcs that increase the size of the largest high-DoS connected component, after which the GA evolutionary operators are resumed. Finally, we include a budget-saturation step to avoid leaving substantial funds unused. Let $\tilde{B}$ denote the residual budget remaining after evaluating a candidate solution $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, and let Δ_a be the cost of upgrading arc a to DoS 6, i.e., $\Delta_a = \pi(\tilde{w}_a^{\mathrm{b}}, 6)$. Whenever $\tilde{B} > 0$, we greedily allocate the remaining budget by repeatedly upgrading the currently lowest-DoS arcs that are still affordable. This post-processing increases budget utilization while preserving feasibility. Although the GA-U does not guarantee global optimality, it provides high-quality solutions in a computationally tractable manner for both formulations as presented in the result section 3.5. The pseudo-code of this algorithm is presented in Algorithm 2.

Algorithm 2 Warm Start Genetic Algorithm (GA-U)

Input: Initial DoS vector $\boldsymbol{w}^{\mathrm{b}}$, budget B
Parameters: population size n_p, max generations n, mutation rate n_m, stagnation limit n_0
Output: Best upgrade plan $\tilde{\boldsymbol{w}}^{\mathrm{b}}$
Fitness: $F(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\|$
Initialize population Ω_0 of size n_p with feasible upgrade plans in Ω (cost $\leq B$)
Evaluate $F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ for all $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega_0$
$\tilde{\boldsymbol{w}}^{\mathrm{b}^\star} \leftarrow \arg\max_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega_0} F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$
$c_{\mathrm{stag}} \leftarrow 0$ ▷ stagnation counter
for $t = 1$ to n **do**
    Select parents from Ω_0
    Generate offspring by crossover
    Mutate each offspring with probability n_m
    Repair or discard infeasible offspring (cost $> B$) to enforce Ω
    Apply budget-saturation post-processing (Algorithm 3)
    Evaluate F for offspring
    Form next population Ω_0 (keep best individuals)
    $\tilde{\boldsymbol{w}}^{\mathrm{b}}_{\mathrm{new}} \leftarrow \arg\max_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega_0} F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$
    if $F(\tilde{\boldsymbol{w}}^{\mathrm{b}}_{\mathrm{new}}) > F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ **then**
        $\tilde{\boldsymbol{w}}^{\mathrm{b}^*} \leftarrow \tilde{\boldsymbol{w}}^{\mathrm{b}}_{\mathrm{new}}$
        $c_{\mathrm{stag}} \leftarrow 0$
    else
        $c_{\mathrm{stag}} \leftarrow c_{\mathrm{stag}} + 1$
    end if
    if $c_{\mathrm{stag}} \geq n_0$ **then**
        Diversify: apply one MLCC-style expansion step to $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}$
        $c_{\mathrm{stag}} \leftarrow 0$
    end if
end for
return $\tilde{\boldsymbol{w}}^{\mathrm{b}^*}$

Algorithm 3 Budget Efficiency

Input: Candidate weights $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, remaining budget $\tilde{B}$
Output: Updated weights $\tilde{w}^{\mathrm{b}}_a$, remaining budget $\tilde{B}$
for $a \in \mathcal{A}^{\mathrm{b}}$ sorted by DoS(a) **do**
    Compute upgrade cost $\Delta_a = \pi(\tilde{w}^{\mathrm{b}}_a, 6)$
    if $\tilde{B} \geq \Delta_a$ **then**
        $\tilde{w}^{\mathrm{b}}_a \leftarrow 6$
        Update budget $\tilde{B} \leftarrow \tilde{B} - \Delta_a$
    end if
    if $\tilde{B} \leq 0$ **then**
        Exit
    end if
end for
Return: Updated $\tilde{w}^{\mathrm{b}}_a$, remaining $\tilde{B}$

Bayesian optimization. Bayesian optimization seeks to maximize (or minimize) a black-box objective F over the feasible set Ω using as few expensive evaluations as possible. It maintains an observation set $\mathcal{D}$ containing all tested upgrade plans and their corresponding objective values. Given $\mathcal{D}$, it fits a probabilistic surrogate model $\mathcal{M}$ (e.g., a Gaussian process)

that approximates F and provides both a mean prediction and an uncertainty estimate. At each iteration, the next candidate plan is chosen by maximizing an acquisition function $\mathcal{A}$, which balances exploitation (high predicted objective) and exploration (high uncertainty). The selected plan is then evaluated with the true objective, appended to $\mathcal{D}$, and the surrogate $\mathcal{M}$ is refit. After a fixed evaluation budget, the best observed plan is returned.

Algorithm 4 Bayesian optimization

Input: initial DoS vector $\boldsymbol{w}^{\mathrm{b}}$; budget B
Parameters: initial samples n_{init}; total evaluations n_{eval}
Feasible set: $\Omega = \{\tilde{\boldsymbol{w}}^{\mathrm{b}} : \mathrm{cost}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \leq B,\ w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6\ \forall a\}$
Objective: $F(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \|\hat{y}_{\boldsymbol{w}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}}\|$
Output: best upgrade plan $\tilde{\boldsymbol{w}}^{\mathrm{b}}$
Sample n_{init} plans $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega$ and evaluate $F(\tilde{\boldsymbol{w}}^{\mathrm{b}})$
$\mathcal{D} \leftarrow \{(\tilde{\boldsymbol{w}}^{\mathrm{b}}, F(\tilde{\boldsymbol{w}}^{\mathrm{b}}))\}$ ▷ observations (plan, objective)
Fit surrogate model $\mathcal{M}$ to $\mathcal{D}$ ▷ probabilistic approximation of F
for $t = n_{\mathrm{init}} + 1$ to n_{eval} **do**
    Select $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star} \leftarrow \arg\max_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega} \mathcal{A}(\tilde{\boldsymbol{w}}^{\mathrm{b}}; \mathcal{M})$ ▷ acquisition
    Evaluate $y^\star \leftarrow F(\tilde{\boldsymbol{w}}^{\mathrm{b}^\star})$
    Update $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}, y^\star)\}$
    Refit/update surrogate model $\mathcal{M}$ using $\mathcal{D}$
end for
return $\tilde{\boldsymbol{w}}^{\mathrm{b}^\star}$ ▷ best observed plan

3.5 Benchmark Instances and Case Studies

Our goal now is to compare the performance of our MLCC algorithm, developed to tackle Problem 3.2, against a straightforward benchmark example where the optimal solution is determined by exhaustively evaluating all possible combinations. Additionally, we will compare it with three solvers to measure the distances to optimality.

Let $B = \alpha \cdot B_{\max}$, be the available budget for optimization, with $B_{\max}$ being the budget required to reach the FS graph (with corresponding prediction $\hat{y}_{FS}$), and $\alpha \in [0, 1]$. The cost per meter to upgrade the infrastructure along an arc is set to reflect different construction costs between levels, following Cerema's cost [110] studies. The cost matrix Π is adapted from Cerema's reported unit construction costs for cycling infrastructure upgrades in France [110].

$$\Pi = \left(\begin{array}{c|cccccc} w_a^{\mathrm{b}} \backslash \tilde{w}_a^{\mathrm{b}} & 1 & 2 & 3 & 4 & 5 & 6 \\ \hline 1 & 0€ & 150€ & 250€ & 300€ & 500€ & 750€ \\ 2 & 0€ & 0€ & 100€ & 150€ & 350€ & 600€ \\ 3 & 0€ & 0€ & 0€ & 50€ & 250€ & 500€ \\ 4 & 0€ & 0€ & 0€ & 0€ & 200€ & 450€ \\ 5 & 0€ & 0€ & 0€ & 0€ & 0€ & 250€ \\ 6 & 0€ & 0€ & 0€ & 0€ & 0€ & 0€ \end{array}\right) \tag{3.43}$$

We select 10 logarithmically spaced values of α between 0.01 and 1 to examine the behavior of our analysis near zero.

$$\alpha_i = 10^{\left(a + i \cdot \frac{b-a}{9}\right)}, \quad i \in \{0, 1, \ldots, 9\}, \tag{3.44}$$

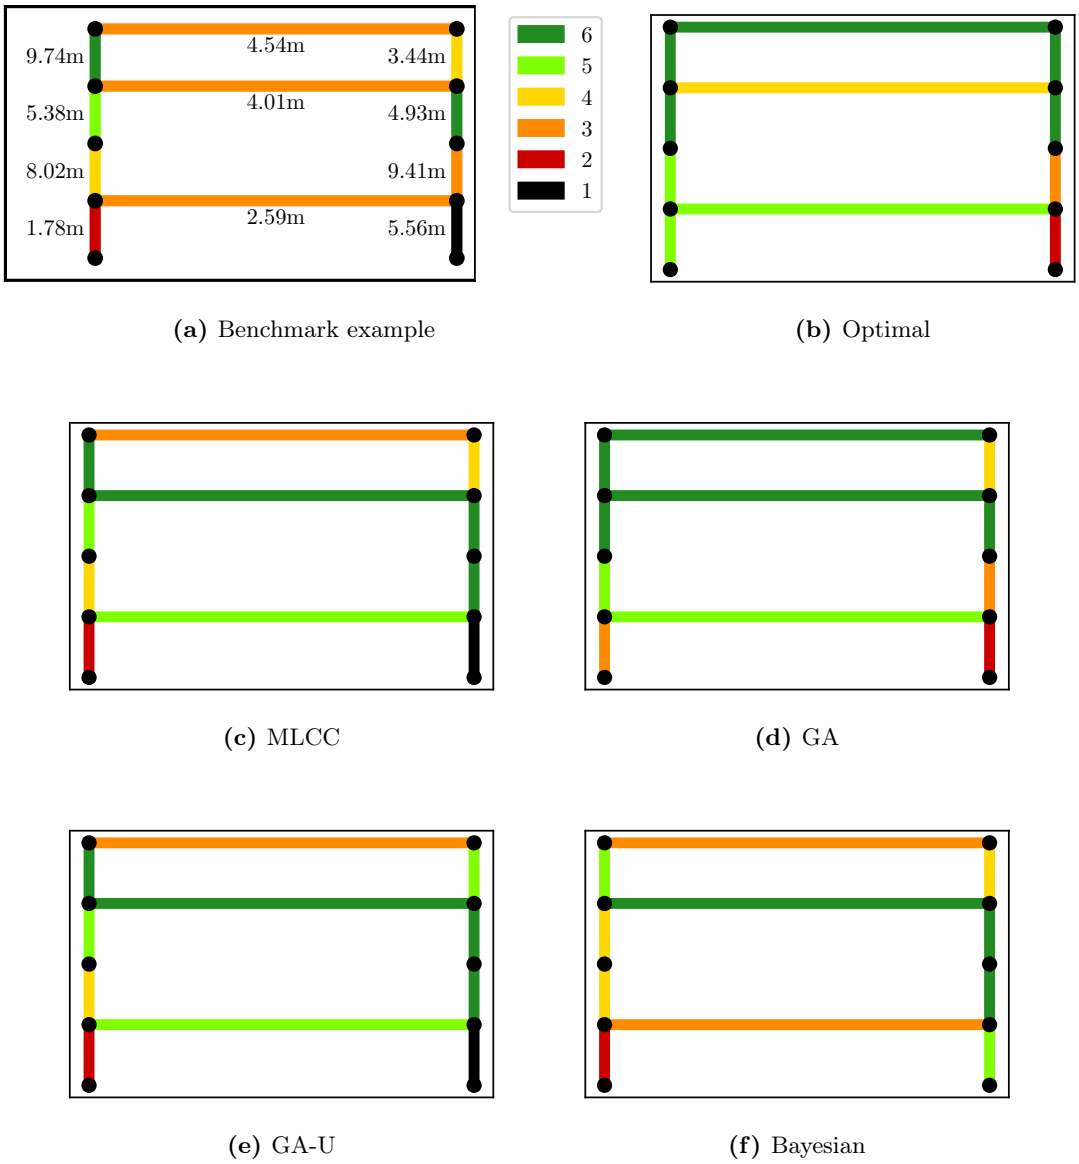

(a) Benchmark example **(b)** Optimal

(c) MLCC **(d)** GA

(e) GA-U **(f)** Bayesian

Figure 3.7: (a) Original configuration, (b) optimal solution configuration, (c) MLCC algorithm configuration, (d) GA configuration, (e) GA-U configuration and (f) Bayesian algorithm configuration: These configurations aim to enhance perceived safety with $i = 7$, corresponding to $\alpha_7 \approx 0.35\%$ of the total budget required for a full network upgrade ($B_{\max}$ =22,010€).

where

$$a = \log_{10}(0.01), \quad b = \log_{10}(1). \tag{3.45}$$

In Fig. 3.7 (a), we display our benchmark example consisting of 10 nodes and 11 arcs. Fig. 3.7 (b)-(f) compare the configurations as explained in Section 3.4 derived from: (1) the optimal solution obtained by evaluating all 138,240 possible combinations via brute force, (2) our MLCC algorithm, (3) a standard genetic algorithm (GA) [111], (4) the GA initialized with the

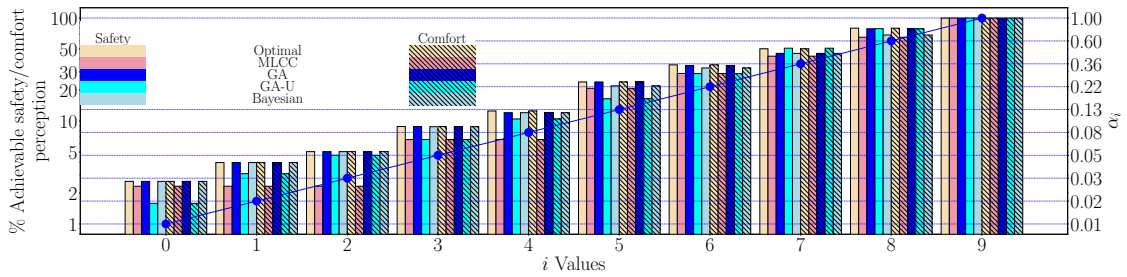

(a)

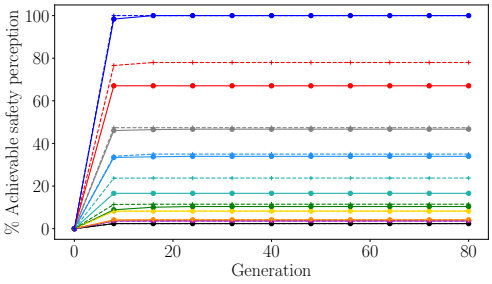

(b)

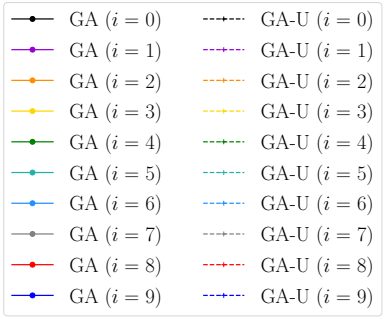

(c)

Figure 3.8: (a) Results of solving Problem 3.2 using the optimal, MLCC, GA, GA-U, and Bayesian algorithms on the benchmark network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) Performance evolution of the GA and GA-U algorithms over generations for safety with (c) providing the legend. $B_{\max} = 22,010$ €.

Algorithm	**Time**									
	$i = 0$	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	$i = 7$	$i = 8$	$i = 9$
Optimal	4.96 s	4.33 s	4.35 s	4.46 s	6.0 s	11.19 s	32.62 s	1.63 m	3.04 m	4.7 m
MLCC	0.014 s	0.011 s	0.011 s	0.011 s	0.011 s	0.011 s	0.011 s	0.014 s	0.012 s	0.011 s
GA	6.7 s	6.99 s	7.28 s	8.39 s	8.46 s	8.72 s	7.58 s	8.91 s	9.89 s	10.6 s
GA-U	7.8 s	8.2 s	8.35 s	8.6 s	7.45 s	7.69 s	9.18 s	9.22 s	10.16 s	10.82 s
Bayesian	1.35 m	1.45 m	1.45 m	1.63 m	1.88 m	1.98 m	1.96 m	1.89 m	1.96 m	2.03 m

Calculation time for the benchmark example

output of the MLCC algorithm (GA-U), and (5) a Bayesian optimization package in Python [112], all designed to approximate the optimal solution within a reasonable time frame. These comparisons are made to approximate the optimal solution for safety using approximately one third of the budget ($\alpha_7 \approx 0.35$). While our MLCC algorithm does not precisely reproduce the optimal solution, it exhibits strong similarities with all other algorithms. Notably, all solutions prioritize maximizing the connectivity of DoS = 6 score segments rather than upgrading isolated links. This is because enhancing a single, unconnected segment to a high DoS = 6 scores would inefficiently allocate the budget without significantly improving perceived safety and comfort across the network, demonstrating the MLCC algorithm's strong performance in delivering a practical and computationally efficient solution for larger networks, while acknowledging that the exact optimality gap remains uncertain. In Fig. 3.8 (a), we illustrate

the improvements in normalized perceived safety and comfort across four different values of α. At α_7, both the optimal solution and the Bayesian algorithm achieve approximately a 29% improvement in perceived safety and comfort, whereas our algorithm achieves around 26%. Although our algorithm does not reach the optimal level, it provides a favorable outcome.

Fig. 3.8 (b) shows the evolution of GA and GA-U over 80 generations, respectively. This figure demonstrates that, in most cases, GA-U approaches the optimal solution more closely than the standard GA. This suggests that using the MLCC results as the initial population for the GA is a beneficial strategy. Finally, bottom of Fig. 3.8 presents a table with the computation times, showing that the MLCC operates significantly faster, making it especially beneficial for large-scale, real-world cycling networks.

One-shot budget (no staging) In Problem 3.2, the budget B is interpreted as a single, static investment envelope allocated once to produce one upgraded configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$. The different budget values explored in Section 3.5 are independent scenarios used to study sensitivity to available resources; they are not consecutive stages over time. Accordingly, this chapter does not model a multi-year roll-out, discounting, or net-present-value (NPV) accounting. Extending the formulation to staged investments with time-indexed decisions and discounted costs/benefits is left for future work.

Sioux Falls and Grenoble Cycling Network

To illustrate a well-known transportation research example, we use the Sioux Falls network (24 nodes, 38 arcs) [113] with randomly assigned DoS values (Fig. 3.10).

As network size grows, computing an optimal solution becomes infeasible due to exponential complexity. For example, evaluating one configuration takes 0.009 seconds, but testing all possibilities would take over a trillion years, highlighting the need for efficient algorithms like MLCC. Fig. 3.11 shows that MLCC performs competitively, with GA-U consistently achieving the best results. Our benchmark analysis in the previous section suggests that the algorithm is relatively close to the optimal solution, and after validating it on the well-known Sioux-Falls benchmark, we now extend our approach to a real case: the Grenoble cycling network. Fig. 3.14 (a) compares algorithm performance for the Grenoble network, showing that MLCC consistently outperforms others. Fig. 3.14 (b) tracks GA and GA-U performance over 80 generations, revealing that GA-U fails to surpass MLCC and ultimately converges to the same solution. Despite parameter adjustments and initial population variations, MLCC consistently yields the best performance. The table at the bottom of Fig. 3.14 demonstrates MLCC's advantage, achieving the best results with significantly lower computation time. This underscores its efficiency in solving large network problems, making it the preferred choice over heuristic methods. Fig. 3.13 visualizes the upgraded Grenoble cycling network using MLCC with α_7. Extending this approach, MLCC is applied to optimize cycling networks across 14 French cities.

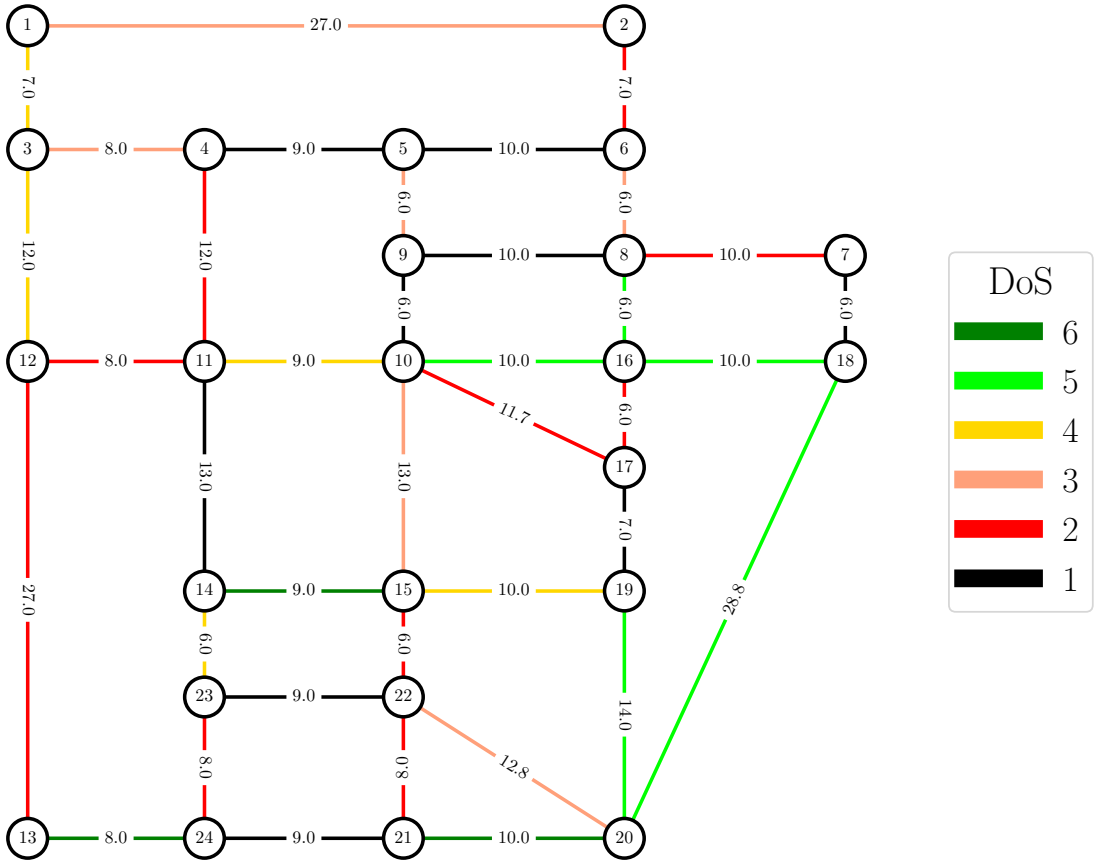

Figure 3.10: Sioux Falls cycling network

Case-Study: Fourteen French Municipalities

In Fig. 3.16, we present the results of the MLCC algorithm, highlighting outcomes for both safety and comfort across various cities. The parameterization specified in Eq.(3.35) leads to notable similarities among the cities. Cities such as Annecy and La Rochelle show significant improvements, approaching nearly 80% for $\alpha_8 \approx 0.6$. This is mainly due to the presence of many disconnected components in their cycling networks. The enhancement of the largest connected component has a particularly substantial impact, with almost complete network coverage leading to these high levels of perceived safety and comfort.

Conversely, cities like Angers, Bordeaux, Dunkerque, Nantes, Rennes, and Tours exhibit relatively minor improvements in safety and comfort. These networks, though containing high-quality infrastructure (high-DoS), are often scattered and sparse, so a significant portion of the budget is spent on connecting these dispersed high-DoS components within the larger connected component, indicating possibly inadequate planning where high-DoS roads

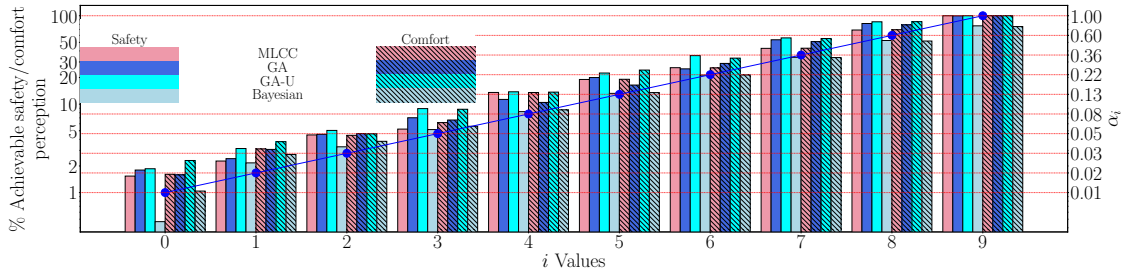

(a)

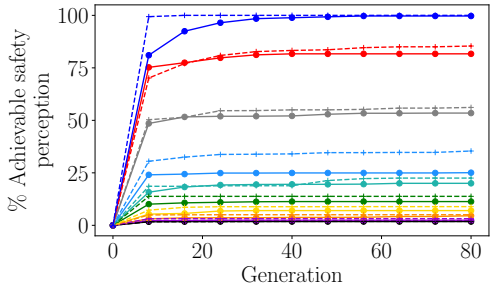

(b)

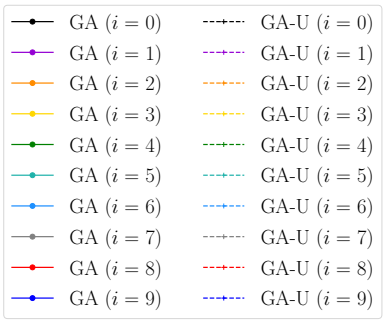

(c)

Figure 3.11: (a) Results of solving Problem 3.2 using the MLCC, GA, GA-U, and Bayesian algorithms on the Sioux Falls cycling network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) show the performance evolution of the GA and GA-U algorithms over generations for safety with (c) providing the legend. And $B_{\max} = 194{,}848$€.

Algorithm	**Time**									
	$i = 0$	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$	$i = 6$	$i = 7$	$i = 8$	$i = 9$
MLCC	0.08 s	0.09 s	0.08 s	0.07 s	0.08 s	0.1 s	0.1 s	0.1 s	0.14 s	0.14 s
GA	25.66 s	13.04 s	12.98 s	27.6 s	14.74 s	17.91 s	29.0 s	20.96 s	23.4 s	40.14 s
GA-U	13.19 s	12.54 s	12.74 s	15.07 s	14.33 s	15.63 s	17.12 s	20.12 s	23.57 s	27.32 s
Bayesian	4.1 m	4.03 m	4.06 m	5.34 m	5.96 m	6.01 m	5.67 m	5.69 m	6.25 m	6.42 m

Calculation time for the Sioux Falls cycling network.

are intermixed with less safe and inconvenient routes.

Cities such as Grenoble, La Rochelle, Le Havre, Lyon, and Paris show considerable improvements of around 5 % with $\alpha_0 = 0.1$, almost because their high-DoS components are relatively close to each other and sufficiently large. Connecting these components to the largest high-DoS connected component results in significant changes, suggesting effective urban cycling planning by city authorities. On the other hand, Rennes and Strasbourg show limited improvements for the smallest values of α_i. Although these cities have a relatively large number of ars, as shown in Table 3.5, they also feature extensive high-DoS connected components that are dispersed among unsafe and uncomfortable roads. This causes the algorithm to initially allocate the budget to upgrading links to connect these components, resulting in

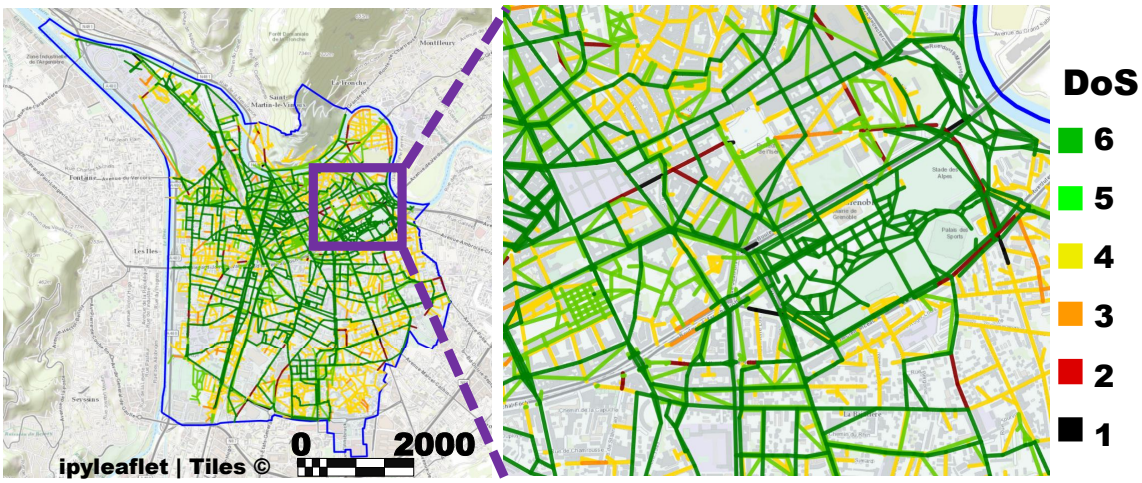

Figure 3.13: Cycling network of Grenoble, France, after being upgraded using MLCC with $\alpha_7 \approx 0.35$.

No	City	$\lvert\mathcal{V}^{\mathrm{b}}\rvert$	$\lvert\mathcal{A}^{\mathrm{b}}\rvert$	Length m	€ $B_{\max}$
1	Angers	7881	10198	546 854 m	307 230 491 €
2	Annecy	4375	5375	301 133 m	170 461 482 €
3	Bordeaux	7190	9850	547 433 m	291 560 346 €
4	Chambery	2266	2544	128 491 m	64 125 336 €
5	Dunkerque	4718	6375	376 446 m	196 435 124 €
6	Grenoble	6773	9017	415 213 m	236 652 715 €
7	La Rochelle	7977	11096	489 185 m	278 834 295 €
8	Le Havre	2609	3043	371 967 m	176 844 796 €
9	Lyon	11857	18418	923 783 m	586 421 676 €
10	Nantes	17194	23529	1 106 161 m	682 956 132 €
11	Paris	33626	54536	2 590 580 m	1 569 036 148 €
12	Rennes	11081	14229	732 691 m	379 182 429 €
13	Strasbourg	11417	16503	649 650 m	247 298 690 €
14	Tours	8723	12380	591 725 m	343 172 657 €

Table 3.5: The cities, the number of nodes, arcs, total amount of meters and the budget required to upgrade all arcs to the highest possible DoS.

minimal immediate improvements.

These results underscore the importance of a linked urban cycling infrastructure, where high quality, connected and continuous cycleways are critical to maximizing perceived safety and comfort. The varying degrees of success between cities reflect the effectiveness of different planning strategies. In particular, cities with more connected and integrated networks in their largest connected component experience superior improvements. Consequently, this analysis

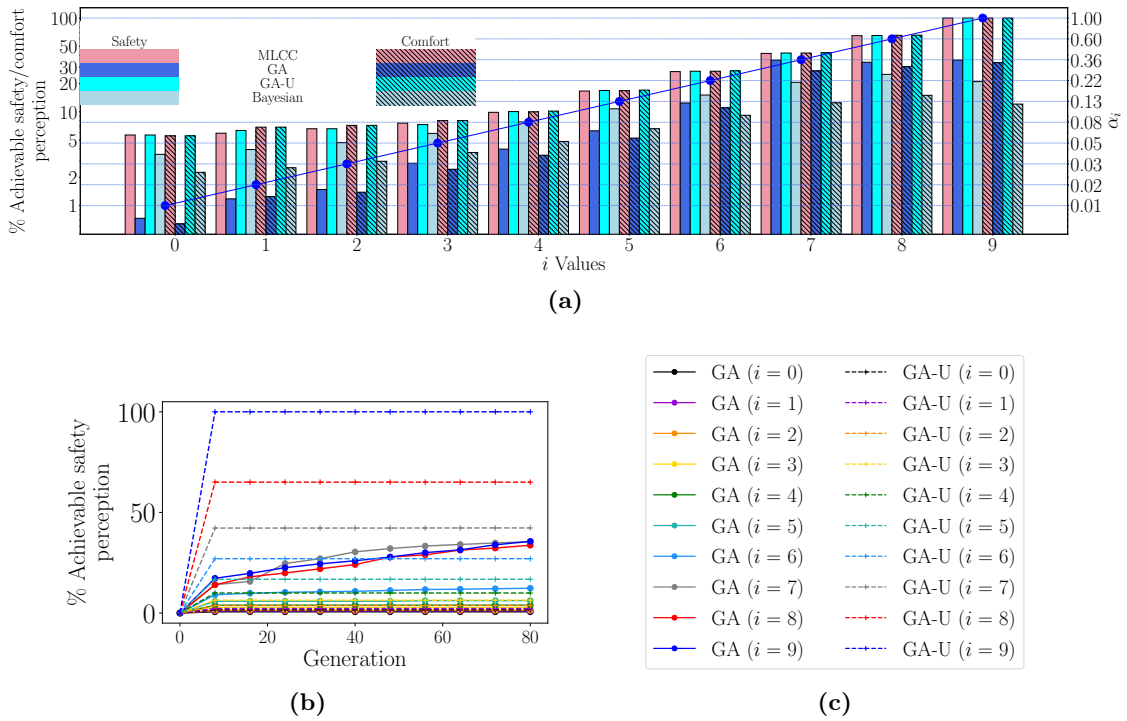

Figure 3.14: (a) Results of solving Problem 3.2 using the MLCC, GA, GA-U, and Bayesian algorithms on the Grenoble cycling network. The percentage improvement is calculated as $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathrm{b}^*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$. (b) Performance evolution of the GA and GA-U algorithms over generations for safety, with (c) providing the legend.

Algorithm	**Time**									
	$i=0$	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$	$i=6$	$i=7$	$i=8$	$i=9$
MLCC	0.48 s	0.48 s	0.53 s	0.62 s	0.82 s	0.9 s	1.28 s	3.4 s	7.88 s	40.0 s
GA	0.46 m	0.45 m	0.47 m	0.46 m	0.51 m	1.62 m	0.34 hr	1.10 hr	1.23 hr	1.20 hr
GA-U	0.49 m	0.53 m	0.99 m	2.74 m	6.82 m	0.28 hr	0.43 hr	1.02 hr	1.34 hr	1.54 hr
Bayesian	1.03 hr	1.01 hr	1.02 hr	1.01 hr	1.34 hr	1.46 hr	1.53 hr	2.48 hr	2.02 hr	2.06 hr

Calculation time for the Grenoble cycling network.

provides valuable insights for urban planners. Ensuring that high-quality cycling infrastructure is widespread and well-connected can greatly enhance the cycling experience, suggesting the need for city authorities to design more strategic networks to prioritize cohesive and integrated cycling routes.

3.6 Concluding Remarks

This chapter developed a graph-based approach to quantify and improve the safety and comfort of cycling networks. We first introduced a set of candidate graph-theory metrics to capture key topological properties of cycling infrastructure. A central methodological step was to

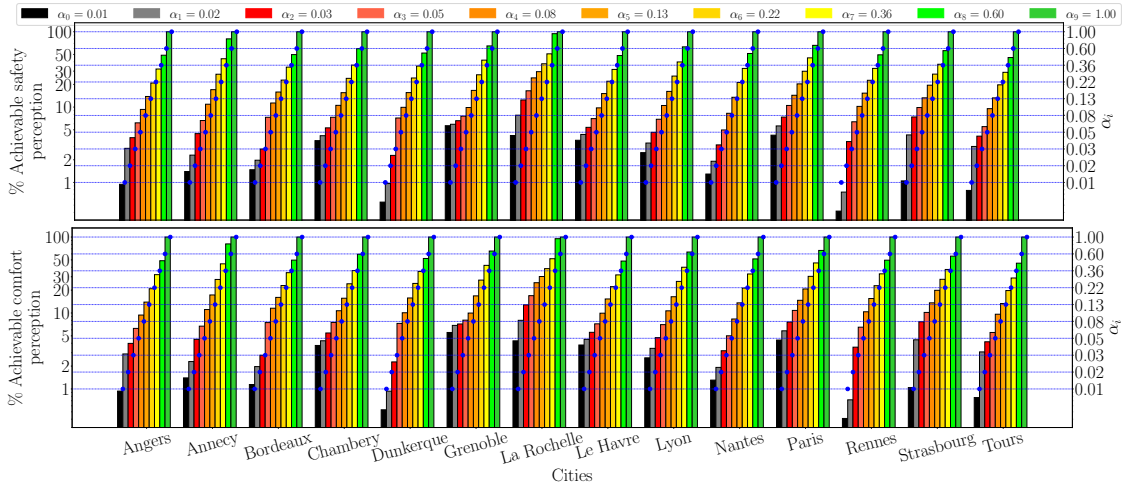

Figure 3.16: Results of the optimization problem (MLCC algorithm) for safety (top) and comfort (bottom) for logarithmically spaced values of α. $\% = |\hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}}} - \hat{y}_{\tilde{\boldsymbol{w}}^{\mathfrak{b}*}}| \times 100/|\hat{y}_{FS} - \hat{y}_{\tilde{\boldsymbol{w}}}|$.

study their monotonic behavior with respect to DoS upgrades, since this property is essential for formulating a meaningful upgrade problem: when an infrastructure segment is improved (higher DoS), the chosen metrics should not behave erratically. This analysis allowed us to retain metrics that are both interpretable in a transport context and suitable for optimization.

Using these validated metrics, we then linked network structure to cyclists' perceived safety and comfort by exploiting survey-based evaluations across French municipalities. The resulting model provides a practical way to approximate how cyclists perceive a network from its topology and its DoS configuration. Building on this model, we formulated a budget-constrained topology optimization problem whose goal is to select upgrades that maximize the predicted safety and comfort level of the network.

Because the resulting problem is combinatorial and highly non-linear, we proposed a reduced-complexity heuristic, the MLCC algorithm, designed to efficiently improve the connectivity of highly protected (high-DoS) corridors, and we compared it with three alternative heuristic strategies. The framework was evaluated on 14 French cities using real infrastructure data and survey scores. Overall, the results show that meaningful gains can often be achieved with limited investment, especially in cities where high-DoS facilities exist but remain fragmented (i.e., isolated segments rather than continuous connected corridors). In such cases, targeted upgrades that connect these segments can substantially improve the global safety/comfort perception predicted by the model.

Despite its strengths, this work has clear limitations. The safety and comfort scores are inferred largely from network-level graph metrics, which can smooth out local hazards such as intersections, crossings, or short conflict points. In addition, the current formulation remains largely topology-driven: it does not explicitly represent cyclist route choice or how demand is distributed across the network, which can result in upgrades that are structurally attractive but affect relatively few trips. Finally, distributional questions are not addressed. The frame-

work does not evaluate whether investments are fairly allocated across socioeconomic areas or aligned with where potential users live and travel; incorporating such considerations would require additional contextual data (e.g., population density, income, and access to opportunities), which lies outside the scope of this chapter.

Future work can build on this foundation in several ways. Prediction quality could be improved by complementing DoS-based topology with features that capture physical effort and routing preferences, such as slope, turning effort (angular change), and cyclist flows. Where available, crash or near-miss records could support calibration and validation against objective safety outcomes, alongside perception-based measures. On the optimization side, embedding the upgrade problem in a flow-aware (e.g., bi-level) setting would allow infrastructure changes to influence route choice and demand, leading to interventions that better reflect actual usage. In addition, a natural extension is to move from the present one-shot budget setting to a staged investment model (multi-period budgets), where the timing of upgrades and discounting (NPV) can affect which projects should be prioritized first. Finally, extending the scope to multi-modal interactions including potential mode shift and incorporating transport-justice objectives would enable a more complete assessment of both the network wide benefits and who gains from the investments. In addition, future work should revisit the restriction to monotone objectives. While monotonicity is valuable for interpretability when DoS upgrades are treated as unambiguous improvements, non-local and potentially non-monotone metrics (e.g., betweenness, eigenvector, or walk-based accessibility) may capture network-wide dependencies that matter for navigation and resilience. More broadly, once route choice and congestion effects are included, improvements can generate system-level trade-offs (in the spirit of Braess-type effects), so exploring alternative objectives including robustness-oriented or equilibrium-aware criteria is an important extension.

Finally, to support cross-mode decision making, the framework should be complemented with a monetized appraisal layer (welfare / benefit cost analysis). Valuing travel-time changes, safety outcomes, and health benefits against investment and external costs would provide an efficiency baseline and allow the cost of alternative design principles (including equity-oriented objectives) to be reported in comparable units, consistent with standard transport appraisal practice.

Chapter 4

Network Design with Flow Redistribution

Contents

4.1 Introduction 71
4.1.1 Summary of Results 72
4.1.2 Related Work 72
4.2 Model Design 73
4.2.1 Cyclist Flow and Exposure 73
4.2.2 Optimization Problem 76
4.3 Case Study 78
4.3.1 Parameters and Algorithm Comparisons 78
4.3.2 Algorithms Comparison 80
4.4 Concluding Remarks 82

4.1 Introduction

In the previous chapter, motivated by evidence that cyclists strongly prefer facilities that are well separated from motorized traffic [2–4], we developed an upgrade framework that models safety and comfort through the DoS of bike lanes. The central idea was that higher DoS values correspond to safer and more comfortable cycling conditions, and that increasing DoS on selected links can improve the overall structure of the cycling network. To guide where upgrades should be made, we relied on graph-theory metrics that predict perceived safety and comfort. These metrics are useful because they quantify whether safe links form a coherent network rather than isolated fragments, and whether cyclists can reach destinations through multiple safe options when disruptions or detours occur. However, a purely topological perspective is not sufficient for infrastructure investment decisions. A link may be structurally important, for instance, it may connect two dense parts of the network, yet still carry limited cyclist demand in practice. Upgrading such a link could therefore have a small effect on the experience of most cyclists, leading to an inefficient allocation of scarce public resources.

This chapter shifts the focus from topology-driven upgrades to flow-aware planning. We explicitly incorporate cyclist flows into the decision process, with the goal of prioritizing upgrades that not only increase safety and comfort, but also benefit a meaningful share of cyclists. To do so, we introduce a safety metric that accounts for cyclist flow and we model how infrastructure improvements change route choice and, consequently, how cyclists redistribute across the network. Beyond proposing this demand-sensitive framework, we also use it to revisit the conclusions of the previous chapter: we compare the upgrade priorities and network-level outcomes obtained under the flow-aware objective with those obtained under the purely topological objective. This comparison suggests that minimizing exposure provides improvement profiles that largely align with the connectivity improvements highlighted by the topological strategy, while offering a more robust basis for prioritization, as it takes into account how cyclists actually use the network. In this way, the chapter connects the structural design of the network with demand-based assessment, maintaining the focus on continuous safe routes and directing investment toward the busiest corridors.

4.1.1 Summary of Results

Applying the proposed flow-aware bi-level framework to the Grenoble cycling network shows that incorporating cyclist flows materially changes which links are prioritized for upgrades compared with purely topology-driven strategies. Under a fixed budget, the optimization concentrates investments on corridors that combine (i) low DoS and high traffic exposure, and (ii) high predicted usage under the all-or-nothing assignment. As a result, the upgraded network achieves a clear reduction in aggregate cyclist exposure and produces a redistribution of flows toward safer facilities, improving the safety experienced by a larger share of cyclists rather than mainly improving structurally central but weakly used links.

In the comparative analysis, GA-U (the warm-start genetic algorithm introduced in the previous chapter) consistently identifies upgrade plans that achieve lower total exposure while remaining budget-feasible. This reflects its ability to exploit intermediate DoS upgrades (e.g., upgrading to DoS 5 rather than always targeting DoS 6) to obtain large exposure reductions per unit cost. In contrast, the baseline MLCC strategy—designed to expand the largest connected component of high-DoS infrastructure primarily prioritizes connectivity of the safest class, which can produce upgrades that are less aligned with observed/estimated demand and therefore less effective at reducing exposure on the corridors most cyclists actually use. Overall, the case study confirms that combining infrastructure quality (DoS), traffic intensity (AADT), and demand-driven flow assignment yields upgrade plans that are both topologically coherent and practically impactful, providing a decision-oriented basis for prioritizing cycling investments.

4.1.2 Related Work

Flow modeling and cycling infrastructure optimization have been studied from multiple perspectives. Bi-level formulations are frequently used to capture the interaction between plan-

ning decisions and traveler responses; for example, [56, 57] optimize network design while accounting for behavioral redistribution of cyclists and impacts on motorized traffic, and [61] evaluates staged network expansions using societal cost–benefit analyses. Complementary work uses traffic assignment and route choice models to estimate bicycle flows under network and level-of-service changes; for instance, [84] applies an assignment approach that accounts for route attributes such as slope and angular change. Graph-theory approaches have also been used to guide infrastructure upgrades by improving connectivity and accessibility properties [59, 60]. These methods offer valuable structural insights, but they may prioritize links based on network position alone, without explicitly accounting for where cyclists actually travel. As a result, topologically attractive interventions can be selected even when they affect only a small share of demand. In contrast, our framework integrates safety and comfort improvements with flow redistribution: the upper level selects upgrades that reduce cyclists' exposure to motorized traffic subject to budget and infrastructure constraints, while the lower level models the resulting cyclist flow patterns. This flow-aware design captures the interdependence between comfort, safety, and usage, and supports investment decisions that maximize practical impact.

4.2 Model Design

In this section, we outline the method for distributing regional mobility demand across individual nodes within each region to enrich the bike origin/destination (OD) data. Next, we introduce the model for predicting cyclist flows, followed by the bi-level optimization problem, which aims to minimize cyclists' exposure to traffic while accounting for the distribution of cycling flows across the network to improve safety and comfort. We then present the two algorithms employed to solve this NP-hard problem. Finally, we present the largest CC metric to evaluate the overall perceived safety and comfort achieved by each algorithm.

4.2.1 Cyclist Flow and Exposure

Using the EMD2020 household travel survey [83], we extract mobility information for the Grenoble area, including the spatial distribution of bicycle trip origins and destinations. As illustrated in Fig. 4.1, the city is partitioned into a set of regions $\mathcal{R} = \{1, \dots, R\}$. For each region $r \in \mathcal{R}$, we denote by $\mathcal{V}^{\mathrm{b}}_r \subseteq \mathcal{V}^{\mathrm{b}}$ the subset of cycling-network nodes located in r. Following [56], we model route choice using an all-or-nothing (AoN) assignment, which is appropriate when congestion effects on bicycle travel times are negligible. Under AoN, the full demand $q^{\mathrm{b}}_{(u,v)}$ is assigned to a single shortest path between u and v. Shortest paths are computed with respect to a cyclist-distance cost that captures relevant topological and environmental characteristics of the cycling infrastructure.

The spatial distribution of bicycle flows is obtained through a procedure analogous to betweenness centrality [114], which determines whether an arc lies on the shortest path between any origin–destination pair and, under the AoN assumption, assigns the corresponding de-

Symbol	Description	Domain	Unit
u, v	Nodes in cycling network	$\mathcal{V}^{\mathrm{b}}$	-
$q^{\mathrm{b}}_{(u,v)}$	Bike demand from u to v	$\mathbb{R}_{\geq 0}$	bikes/hr
$p^{\mathrm{b}}_{(u,v)}$	Shortest path between u and v	$\mathcal{A}^{\mathrm{b}}$	-
d^{b}_a	Arc a length/distance	$\mathbb{R}_{>0}$	m
$\measuredangle^{\mathrm{b}}_a$	Angular change at a	$\mathbb{R}_{\geq 0}$	deg
s^{b}_a	Slope factor of arc a	$\mathbb{R}_{\geq 0}$	-
f^{b}_a	Traffic factor on arc a	$\mathbb{R}_{\geq 0}$	-
$\bar{d}^{\mathrm{b}}_{(u,v)}$	Cyclist distance u-v	$\mathbb{R}_{\geq 0}$	m
ϕ^{b}_a	Flow on arc a	$\mathbb{R}_{\geq 0}$	bikes/hr
$\mathbb{I}_{a \in p^{\mathrm{b}}_{(u,v)}}$	arc a is part of shortest path $p^{\mathrm{b}}_{(u,v)}$?	$\{0, 1\}$	-
ρ, κ, ψ	Calibration parameters	$\mathbb{R}$	-
$\mathrm{AADT}(a)$	AADT adjacent to a	$\mathbb{R}_{\geq 0}$	veh/day
w^{b}_a	Original weight of a	$\mathbb{R}_{\geq 0}$	-
$\tilde{w}^{\mathrm{b}}_a$	Candidate weight for a	$\mathbb{R}_{\geq 0}$	-
$\pi(w^{\mathrm{b}}_a, \tilde{w}^{\mathrm{b}}_a)$	Upgrade cost for a	$\mathbb{R}_{\geq 0}$	€
$l_{\mathcal{G}^{\mathrm{b}}}$	Largest connected component metric	$[0, 1]$	-

Table 4.1: Notation summary of the Chapter 4.

mand to that path. For an arc $a \in \mathcal{A}^{\mathrm{b}}$ belonging to the shortest path $p^{\mathrm{b}}_{(u,v)} = (a_1, \ldots, a_k)$, d^{b}_a denotes its geographical length; $\measuredangle^{\mathrm{b}}_a$ the angular change (in degrees) if it is the first or last arc in the path, or at junctions between consecutive arcs; s^{b}_a the slope factor based on the terrain slope m^{b}_a; and f^{b}_a the traffic factor, representing the influence of adjacent car traffic on perceived cycling comfort and safety.

The cyclist-specific distance $\bar{d}^{\mathrm{b}}_{(u,v)}$ between u and v, as defined in [84], is then expressed as:

$$\bar{d}^{\mathrm{b}}_{(u,v)} = \sum_{a \in p^{\mathrm{b}}_{(u,v)}} \left(d^{\mathrm{b}}_a \cdot s^{\mathrm{b}}_a \cdot (f^{\mathrm{b}}_a)^{\rho} \right) + \sum_{a \in p^{\mathrm{b}}_{(u,v)}} \measuredangle^{\mathrm{b}}_a \cdot \kappa, \tag{4.1}$$

where $\rho > 0$ and $\kappa > 0$ are calibration parameters. The term f^{b}_a represents the traffic exposure factor on arc a and is defined as:

$$f^{\mathrm{b}}_a = \begin{cases} \psi \cdot \exp\left(\mathrm{AADT}(a)\right), & \text{if } w^{\mathrm{b}}_a < 5, \\ \psi \cdot 2.03, & \text{if } w^{\mathrm{b}}_a = 5, \\ \psi, & \text{if } w^{\mathrm{b}}_a = 6, \end{cases} \tag{4.2}$$

where ψ is a scaling factor, w^{b}_a denotes the DoS of bike lane a, and $\mathrm{AADT}(a)$ is the annual average daily traffic on adjacent car lanes that in practice is normalized before being used in the exponential. According to [5], the effect of AADT becomes negligible for fully separated cycling infrastructure (DoS = 5 or 6). The factor 2.03 reflects the more than twofold increase in crash risk in pedestrian-shared spaces compared to fully separated facilities [85]. For convenience,

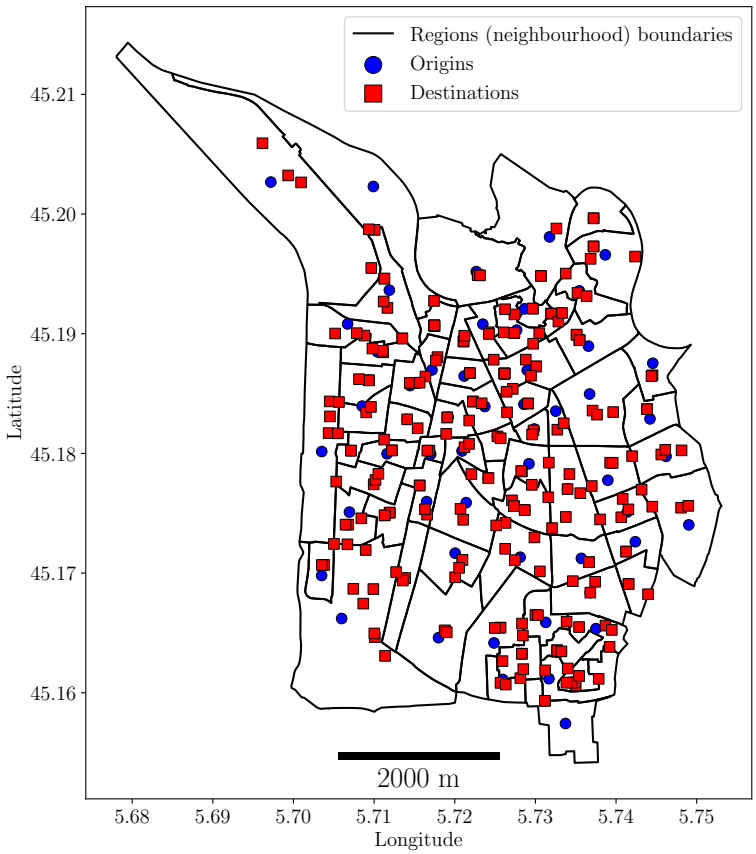

Figure 4.1: Layout of Grenoble city, origin/destinations nodes in regions $\mathcal{R}$.

we also define the cyclist-specific distance contributed by a single arc $a \in p^{\mathrm{b}}_{(u,v)}$ as

$$\bar{d}^{\mathrm{b}}_a = d^{\mathrm{b}}_a \cdot s^{\mathrm{b}}_a \cdot (f^{\mathrm{b}}_a)^{\rho} + \measuredangle^{\mathrm{b}}_a \cdot \kappa. \tag{4.3}$$

Unlike traditional car traffic assignment, where link flows are obtained from solving a user equilibrium problem to reflect congestion and travel-time feedback, the cyclist flow is determined directly from the AoN assignment using the shortest-path structure induced by $\bar{d}^{\mathrm{b}}_{(u,v)}$. This approach assumes that cyclists always choose the most attractive route in terms of perceived distance and safety, without mutual interaction effects, specifically, let $\mathbb{I}_{a \in p^{\mathrm{b}}_{(u,v)}}$ be a binary indicator equal to 1 if arc a belongs to path $p^{\mathrm{b}}_{(u,v)}$, and 0 otherwise. The total cyclist flow assigned to arc a, denoted ϕ^{b}_a, is thus obtained by aggregating all assigned paths:

$$\phi^{\mathrm{b}}_a = \sum_{(u,v) \in \mathcal{V}^{\mathrm{b}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{b}}_{\mathcal{D}}} q^{\mathrm{b}}_{(u,v)} \cdot \mathbb{I}_{a \in p^{\mathrm{b}}_{(u,v)}} \geq 0. \tag{4.4}$$

This ensures that the total flow on each arc is the sum of all demands associated with trips traversing it, reflecting the behavior of the betweenness centrality and maintaining network consistency due to the symmetry of the demand matrix. Finally, the cyclist exposure on a given link $a \in \mathcal{A}^{\mathrm{b}}$ is computed as the product of the cyclist flow, the traffic exposure factor,

and the length of a

$$\mathcal{E} = \sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}} \cdot f_a^{\mathrm{b}} \cdot d_a^{\mathrm{b}}, \tag{4.5}$$

this means that exposure is higher on links where more cyclists are present, traffic conditions are more intense, and the segment covers a greater distance.

4.2.2 Optimization Problem

Taking a utilitarian perspective, we aim to reduce cyclists' average exposure to vehicular traffic while accounting for their flow distribution across the network. By prioritizing upgrades that improve the DoS within a fixed budget, we seek to maximize the average safety and comfort of cyclists, directing flows toward safer routes and reducing high-traffic exposure for all network users.

Each upgrade to the cycling infrastructure, such as increasing a road segment's DoS, incurs a cost per meter. As in Chapter 3 DoS upgrades are discrete and must satisfy a fixed budget B. We define the following notations:

- $\boldsymbol{w}^{\mathrm{b}} = \left[w_1^{\mathrm{b}}, \ldots, w_a^{\mathrm{b}}, \ldots, w_{|\mathcal{A}|}^{\mathrm{b}}\right]$: the original arc weights of the graph $\mathcal{G}$;
- $\tilde{\boldsymbol{w}}^{\mathrm{b}} = \left[\tilde{w}_1^{\mathrm{b}}, \ldots, \tilde{w}_a^{\mathrm{b}}, \ldots, \tilde{w}_{|\mathcal{A}|}^{\mathrm{b}}\right] \in \Omega$: the candidate weights to be optimized, whereby

$$\Omega = \left\{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \{1, \ldots, 6\}^{|\mathcal{A}^{\mathrm{b}}|} : w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6,\ \forall a \in \mathcal{A}^{\mathrm{b}}\right\},$$

Additionally, we define the following variables, all normalized between their respective minimum and maximum values:

Symbol	Description	Domain	Unit
$\phi_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Flow on arc $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	bikes/hr
d_a^{b}	Length of $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	m
$\bar{d}_a^{\mathrm{b}}$	Cyclist distance of $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	m
$\mathcal{E}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Cyclist exposure on $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	-
$\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$	Upgrade cost for $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	€

Table 4.2: Abbreviated notation for variables in cycling network model, writing $(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ next to a variable indicates that it is a function of the candidate configuration for bike lane improvements.

The optimization problem has two interconnected levels:

Upper-Level Problem ($\mathbb{U}_{\mathrm{pper}}$): This problem represents the strategic infrastructure planning goal. Here, we aim to minimize cyclist exposure to car traffic by choosing optimal

upgrades to the DoS across the network, while staying within the budget B. The objective function minimizes the total exposure defined as the product of arc length, traffic factor, and cyclist flow. This ensures that upgrades prioritize long arcs with high traffic factors and flows, as they present the greatest risk of exposure to cyclists.

Lower-Level Problem ($\mathbb{L}_{\text{ower}}$): This problem represents the flow distribution, which models how cyclists will travel across the network once the upgrades have been made. The objective is to determine the optimal distribution of cyclist flows across the network by assigning each OD pair to the path that minimizes cyclist distance by an all-or-nothing assignment as in Section 4.2.1. This distribution depends on the cyclist-specific distance metric, which reflects the comfort and safety of each route.

Because changes in the DoS at the upper level influence the cyclist distance (and thus the flows) at the lower level, the two levels are nested. This means that every decision made at the upper level affects the solution of the lower level. We now optimize a flow-weighted exposure objective instead of topology metrics as follows:

Problem 4.1 (Bi-level network-upgrade optimization problem).

$$\mathbb{U}_{\text{pper}}:$$

$$\tilde{\boldsymbol{w}}^{\text{b}\star} = \arg\min_{\tilde{\boldsymbol{w}}^{\text{b}}\in\Omega} \sum_{a\in\mathcal{A}^{\text{b}}} \phi_a^{b\star}(\tilde{\boldsymbol{w}}^{\text{b}}) \cdot f_a^{\text{b}}(\tilde{\boldsymbol{w}}^{\text{b}}) \cdot d_a^{\text{b}} \tag{4.6}$$

$$\text{s.t.} \sum_{a\in\mathcal{A}^{\text{b}}} d_a^{\text{b}} \cdot \pi(w_a^{\text{b}}, \tilde{w}_a^{\text{b}}) \leq B, \tag{4.7}$$

$$\mathbb{L}_{\text{ower}}:$$

$$\boldsymbol{\phi}^{b\star}(\tilde{\boldsymbol{w}}^{\text{b}}) = \arg\min_{\phi^{\text{b}}\in\mathbb{R}_{\geq 0}^{|\mathcal{A}^{\text{b}}|}} \sum_{a\in\mathcal{A}^{\text{b}}} \bar{d}_a^{\text{b}}(\tilde{\boldsymbol{w}}^{\text{b}}) \cdot \phi_a^{\text{b}}, \tag{4.8}$$

$$\text{s.t. } \phi_a^{\text{b}} = \sum_{(u,v)\in\mathcal{V}_{\mathcal{O}}^{\text{b}}\times\mathcal{V}_{\mathcal{D}}^{\text{b}}} q_{(u,v)}^{\text{b}} \cdot \mathbb{I}_{a\in p_{(u,v)}^{\text{b}}} \geq 0,\ \forall a \in \mathcal{A}^{\text{b}}.$$

The bi-level optimization Problem 4.1 is NP-hard. In particular, it contains a knapsack-type selection problem as a special case: even if the lower-level flow assignment is held fixed, deciding which arc upgrades to apply under a budget constraint requires a combinatorial choice among discrete upgrade actions. In the full bi-level setting, this combinatorial structure is coupled with flow redistribution, since each upgrade changes the shortest-path costs and therefore the all-or-nothing assignment, which further increases the computational burden on large networks.

Remark on monotonicity of exposure. Unlike the topology-based indicators discussed in the previous chapter, the exposure objective in Problem 4.1 is not guaranteed to be monotone

under DoS upgrades when route choice is endogenous. If the lower-level flows were held fixed, then increasing DoS on any arc would weakly decrease exposure, since $f_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ is non-increasing in the DoS level and all other terms in (4.6) remain unchanged. In the bi-level setting, however, upgrades modify the cyclist distance $\bar{d}_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ and can therefore change the AoN shortest paths, leading to discontinuous flow reassignments. As a result, the product $\phi_a^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot f_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \bar{d}_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ may increase on some arcs even if f_a^{b} decreases everywhere, because flows can concentrate on particular corridors or shift toward routes that remain exposed. For this reason, exposure is best interpreted as a policy outcome that captures network and behavioral feedback, rather than as a monotone graph metric. In chapters 6, this non-monotonicity is further amplified when demand responds through modal shift.

4.3 Case Study

In this section, we present the Grenoble (France) case study and compare the upgrade plans produced by GA-U and MLCC. The GA-U method is the genetic algorithm introduced in Chapter 3.4; throughout this chapter, it is used with the objective function of Problem 4.1, so that its search is driven by cyclist-exposure minimization $\mathcal{E}$ under the budget constraint. We then evaluate the resulting solutions not only in terms of exposure, but also in terms of perceived safety and comfort, by computing the largest connected component metric $l_{\mathcal{G}}^{\mathrm{b}}$ for the upgraded networks returned by both algorithms.

4.3.1 Parameters and Algorithm Comparisons

To solve the all-or-nothing traffic assignment problem and determine cyclist flows, as described in Section 4.2.1, we use parameters that are better suited for the model, as suggested by [84]. Specifically, we set $\rho = 0.04$, $\kappa = 0.14$ and $\psi = 0.84$. The slope factor s_a^{b} is determined in Table 4.3, where downhill routes are considered as the reciprocal of the slope factor when traversing in the opposite direction.

Slope (m_a^{b}) %	Uphill Factor (s_a^{b})	Downhill Factor (s_a^{b})
$m_a^{\mathrm{b}} < 2$	1.0	1.0
$2 \leq m_a^{\mathrm{b}} < 4$	1.371	$1/1.371 \approx 0.729$
$4 \leq m_a^{\mathrm{b}} < 6$	2.203	$1/2.203 \approx 0.454$
$m_a^{\mathrm{b}} \geq 6$	4.239	$1/4.239 \approx 0.236$

Table 4.3: Slope factors for cycling distance calculation

Let $B = \alpha \cdot B_{\max}$, be the available budget for optimization, with $B_{\max}$ being the budget required to reach the FS graph, and $\alpha \in [0, 1]$, for this analysis, we assume $\alpha = \frac{1}{2}$, representing 50% of the maximum budget. The cost per meter to upgrade the infrastructure along an edge is set to reflect different construction costs between levels, following Cerema's cost [110] studies as presented in Eq. (3.43).

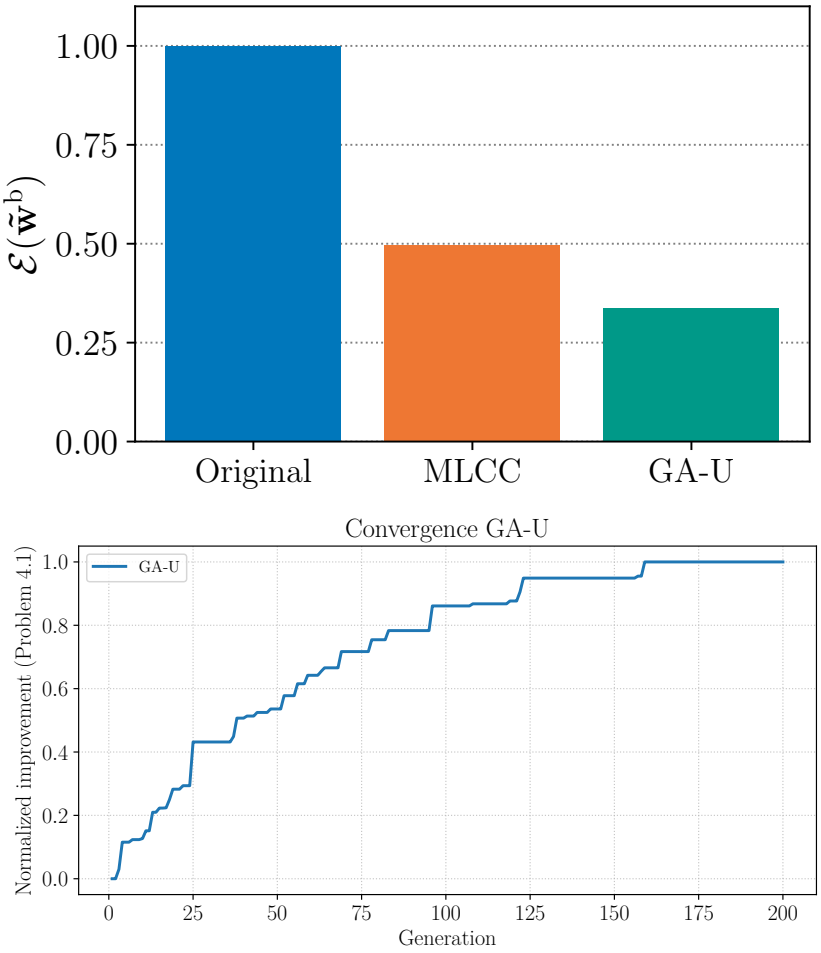

Figure 4.2: On the top is the initial normalized traffic exposure and reduced exposure after optimization using the GA-U and MLCC algorithms for Problem 4.1 with a budget constraint of 50%, i.e., $\alpha = \frac{1}{2}$. And on the bottom the convergence analysis of GA-U for the same budget.

To solve Problem 4.1, we apply the GA-U and MLCC algorithms. Fig. 4.2 compares the original cyclist exposure to traffic of 0.42 with the optimized values of 0.17 and 0.20 obtained with GA-U and MLCC, respectively. The convergence speed suggests that the largest improvements are achieved in the first stages of the search. This behavior is consistent with the structure of the problem: at the beginning of the optimization, the cycling network still contains several highly exposed segments for which a small number of interventions can produce substantial reductions in aggregate exposure. As these high-impact modifications are progressively selected, the remaining feasible upgrades concern parts of the network that are either already improved or contribute less to the objective. The algorithm therefore continues to improve the solution, but the marginal gain per generation becomes smaller. In that sense, the slower progress observed in later generations does not indicate a loss of search efficiency, but rather that the optimization is entering a regime in which most of the major improvements have already been captured. These results indicate that GA-U achieves a larger reduction in

exposure than MLCC, despite providing only a locally optimal solution. Fig. 4.4 reports the percentage distribution of bike-lane length across DoS categories for the original network and for the two upgrade plans. While MLCC concentrates a larger share of the network in the highest category (DoS 6), GA-U assigns a substantially larger share to DoS 5 (approximately five times higher than MLCC). This strategy reflects a more budget-efficient allocation: it shifts a large fraction of infrastructure into a high-separation category that remains cheaper than the maximum DoS level, yet still delivers a marked reduction in cyclist exposure.

Although MLCC does not explicitly consider flow, it still indirectly lowers traffic exposure due to the all-or-nothing traffic assignment model, which assumes cyclists follow the shortest path based on the cyclist distance metric Eq. (4.1). This metric heavily depends on the DoS of each lane, and MLCC indirectly improves OD connectivity with high DoS lanes, providing safe and comfortable paths.

4.3.2 Algorithms Comparison

Fig. 4.3 further illustrates the DoS configuration and the resulting cyclist flow patterns before and after optimization. In this chapter, the total bicycle OD demand is held fixed (no induced demand or destination changes), but route choice is endogenous: cyclists are reassigned through an all-or-nothing shortest-path model based on the generalized cycling cost, where higher DoS reduces the traffic-related risk factor and therefore lowers the perceived cost of using upgraded links. Consequently, the post-upgrade flows shift toward links and corridors that become comparatively more attractive under this exposure-based proxy.

The GA-U solution appears "greener" mainly because, under the budget constraint, it exploits intermediate upgrades (notably to DoS 5) on many links that lie on high-use corridors. These upgrades still yield substantial reductions in the traffic factor f_a^{b} and thus in flow-weighted exposure, while being cheaper than systematically upgrading to DoS 6 as in MLCC. In other words, GA-U achieves larger exposure reductions per unit cost by spreading protection across more of the routes that cyclists actually use, whereas MLCC prioritizes maximizing connectivity of the very safest class, which can concentrate spending on fewer, fully protected links.

Finally, we emphasize that these routing responses are not estimated from local revealed-preference data in this thesis. They should be interpreted as a simplified behavioral mechanism consistent with established evidence that cyclists prefer greater separation from motor traffic. A richer empirical calibration of preferences (e.g., a discrete choice model for route selection and/or induced demand) is an important direction for future work.

Now, we use the largest connected component metric in the different optimization outcomes to asses overall perceived safety and comfort, we visualize the original levels of perceived safety and comfort compared to those achieved by the GA-U and MLCC algorithms. Fig. 4.5 reports the values of the connectivity indicator l_G^{b}, which Chapter 3.3.3 identified as the most informative predictor of perceived safety and comfort. Although the optimization in this chapter targets cyclist exposure $\mathcal{E}$ rather than perceived comfort directly, the upgraded solutions still

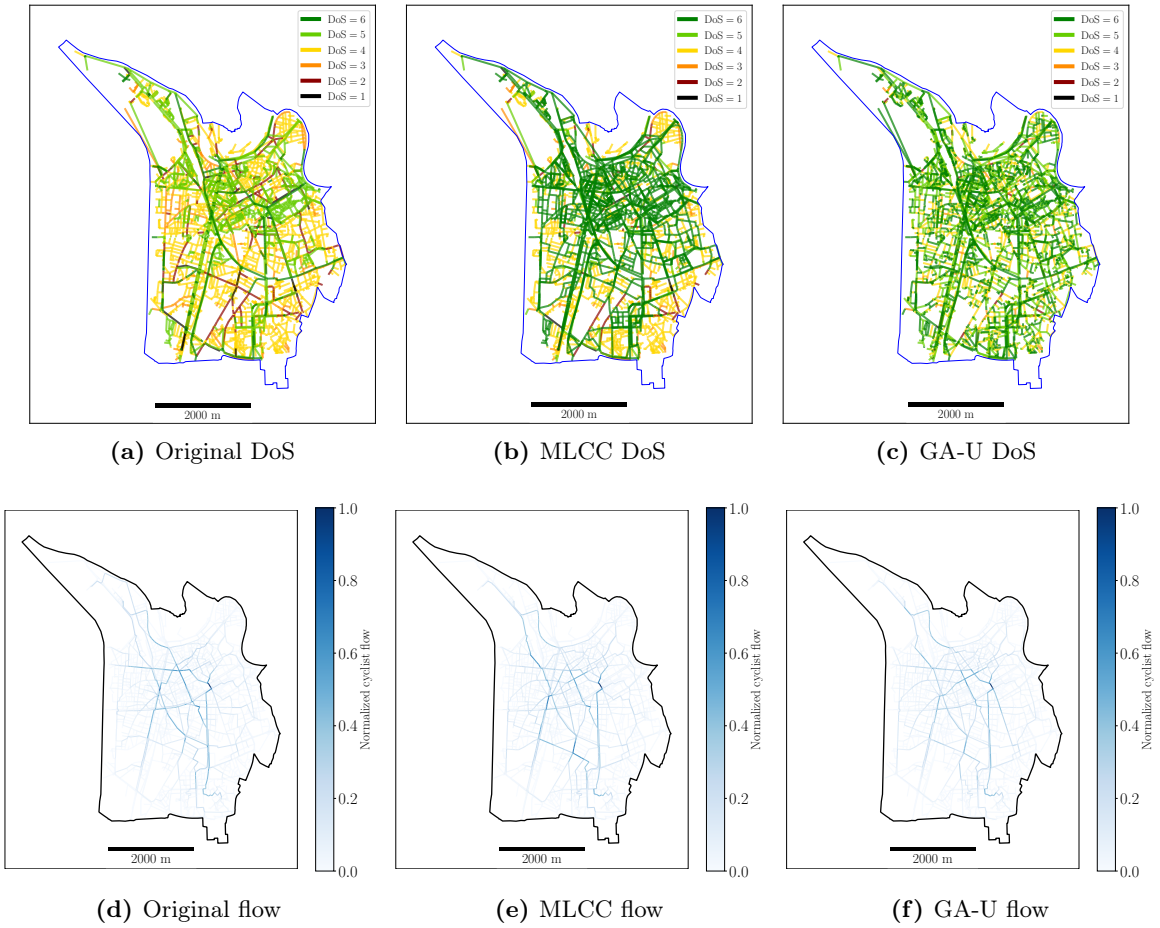

(a) Original DoS **(b)** MLCC DoS **(c)** GA-U DoS

(d) Original flow **(e)** MLCC flow **(f)** GA-U flow

Figure 4.3: The upper row shows the DoS configuration, and the lower row shows the flow mapping. The optimization algorithms operate with 50% of the total budget needed for a complete network upgrade.

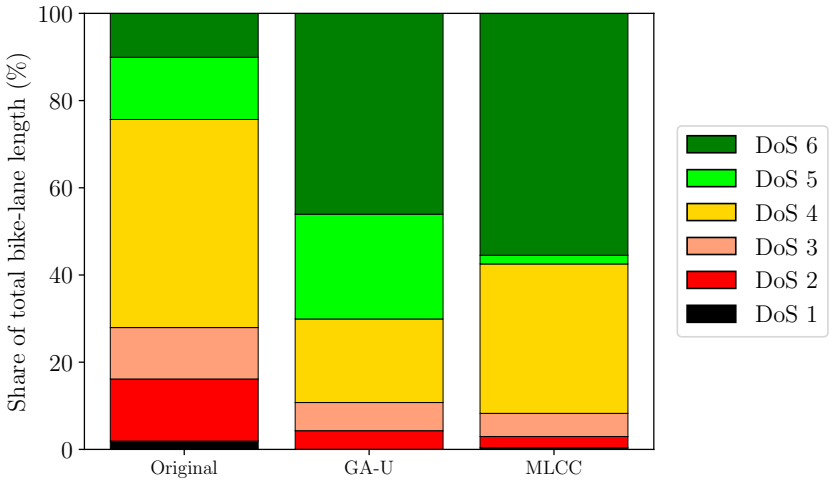

Figure 4.4: Distribution of bike-lane length by DoS (normalized)

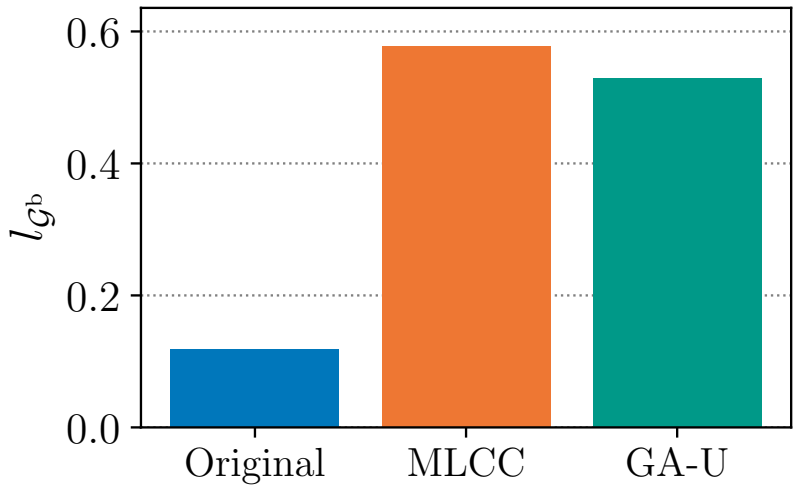

Figure 4.5: Global measure of perceived safety and comfort.

yield a substantial increase in $l^b_{\mathcal{G}}$ relative to the original network (from 0.12 to approximately 0.54 with GA-U and 0.58 with MLCC).

This comparison clarifies the relationship between the two objectives. Minimizing exposure encourages investments on links that matter most for cyclists' experienced risk, which can also strengthen the safe backbone of the network, hence the consistent rise in $l^b_{\mathcal{G}}$. However, the alignment is not perfect: MLCC is designed to maximize the size of the highly protected connected component and therefore achieves the highest $l^b_{\mathcal{G}}$, whereas GA-U allocates part of the budget to high-impact upgrades that may remain isolated from the main high-DoS subnetwork. As a result, GA-U can achieve large exposure reductions while producing a smaller gain in the perceived-safety proxy. Overall, these results support the use of exposure as a practical objective for safety-oriented planning, while also showing that explicitly promoting connectivity of high-DoS facilities can further enhance perceived safety and comfort.

Despite not being explicitly designed to maximize overall perceived safety and comfort, GA-U performs effectively, and both algorithms significantly improve overall perceived safety and comfort. Future work will focus on further optimizing and enhancing overall perceived safety and comfort while accounting for flow distribution and additional factors, such as potential cyclist congestion. Additionally, we aim to develop an algorithm that can approach the global optimum more closely.

4.4 Concluding Remarks

This chapter presented a budget-constrained framework for upgrading cycling infrastructure with the aim of improving cycling conditions at the network level. Compared with the previous chapter, the optimization criterion shifts from a purely topological proxy of perceived safety

and comfort to a flow-aware exposure objective: building on a six-level DoS classification, upgrades are modeled as decisions that reduce cyclists' exposure to motor traffic and, through route-choice effects, influence how cyclists redistribute across the network. While this change of objective reflects a different modeling perspective, we also assess how these exposure-driven upgrade plans translate into perceived safety and comfort by computing the largest connected component metric l_G^{b} on the upgraded networks. The results show that minimizing exposure is generally aligned with improving this perception-oriented indicator, with the additional advantage that exposure explicitly accounts for where cyclists actually travel and therefore prioritizes upgrades that benefit a larger share of users.

To solve the resulting combinatorial problem in practice, we use the genetic algorithm whose initial population is warm-started with MLCC (GA-U), and we compare it with the MLCC strategy that strengthens the continuity of the safest infrastructure. The Grenoble case study shows that targeted upgrades can substantially reduce traffic exposure and promote a more favorable redistribution of bicycle flows toward safer facilities. Overall, the proposed approach provides a practical and scalable tool for prioritizing cycling investments in a way that reflects both infrastructure quality and network usage, while remaining consistent with perception-oriented network structure indicators.

While these results highlight the potential of DoS-based upgrades to improve cycling conditions, they do not yet capture an important system-level trade-off: achieving higher DoS often requires reallocating road space (e.g., removing a general-traffic lane or parking), which can reduce vehicle capacity on the corresponding road segments. In dense urban areas, such capacity reductions may increase car congestion and total travel time, creating impacts that matter for users, planners, and public acceptance. Therefore, the next chapter extends the analysis to a multi-modal setting by explicitly modeling the interaction between cycling upgrades and automobile traffic performance. This provides a first, simplified framework to evaluate cycling-safety benefits alongside potential car-travel-time costs, and to support more balanced infrastructure decisions.

Chapter 5

Network Design with Flow Redistribution in a Multi-Modal Framework

Contents

5.1 Introduction 85
5.1.1 Summary of Results 86
5.1.2 Related Work 86
5.2 Model Design 87
5.2.1 Graph Representation 88
5.2.2 Cars Travel Time 89
5.2.3 Interaction Between Cars and New Bike Lanes 90
5.2.4 Optimization Problem 91
5.3 Case Study 93
5.3.1 Problem Parameters and Optimal Solution 93
5.3.2 Impact on Car Routing Choice and Travel Time 95
5.4 Concluding Remarks 102

5.1 Introduction

In the previous chapters, we focused on upgrading bicycle infrastructure to improve cyclists' safety and comfort, motivated by the well-established evidence that cyclists prefer facilities that are physically separated from motorized traffic [2, 4, 44]. However, such upgrades are rarely implemented in isolation: creating protected or segregated bike lanes often requires reallocating existing road space, which can reduce car-lane capacity and potentially increase congestion and travel times for motor vehicles. This introduces an important planning tension for cities operating under limited budgets: how to prioritize cycling safety improvements while keeping the impact on automobile traffic within acceptable limits.

To address this challenge, we propose a bi-level planning framework that explicitly captures the interaction between cycling upgrades and traffic dynamics. At the upper level, a planner

selects DoS upgrades to reduce cyclists' exposure to traffic while controlling the deterioration of car travel times induced by capacity changes. At the lower level, route choices are modeled through assignment: cyclists choose paths based on a cyclist-specific generalized cost that accounts for factors such as slope, turning angles, and traffic conditions along bike facilities, while car travel times are evaluated through a traffic model consistent with the modified road capacity. This structure enables the evaluation of upgrade plans that are both safety-improving for cyclists and operationally acceptable for motorized traffic.

5.1.1 Summary of Results

The proposed framework highlights that safety-oriented cycling upgrades can be achieved without drastically degrading automobile performance when planning decisions explicitly account for traffic dynamics. In the Grenoble case study, the bi-level model uses upgrade budgets that substantially reduce cyclists' exposure to motor traffic by shifting investment toward higher DoS segments on corridors that carry meaningful bicycle flows. The lower-level bicycle assignment shows that these upgrades not only improve conditions on the upgraded links, but also induce a redistribution of cyclist routes toward safer facilities, leading to a more balanced and safety-oriented flow pattern across the network.

On the motor-traffic side, comparing fixed-route evaluations with post-upgrade user-equilibrium assignment indicates that the car travel-time impact of the selected upgrades remains limited: most of the safety gains for cyclists can be obtained with only marginal increases in total car travel time, even when accounting for congestion feedback after capacity changes. Overall, the results suggest that the best-performing upgrade plans are those that target segments with high cyclist demand while avoiding capacity reductions on already critical car corridors, thereby achieving a favorable trade-off between cyclist safety improvements and multi-modal operational feasibility.

5.1.2 Related Work

The interaction between bicycle network design and automobile traffic has been studied through several modeling and optimization approaches. A common line of work formulates network design problems that improve cycling connectivity while limiting adverse effects on car travel time, often through bi-level optimization in which infrastructure decisions are coupled with user responses [56, 57]. More recent planning frameworks explicitly explore trade-offs in lane reallocation by balancing bicycle and car accessibility, typically using multi-criteria optimization to quantify compromises between modes [62]. Complementary approaches prioritize upgrades on corridors with high cycling potential while imposing constraints on automobile performance, aiming to ensure that cycling improvements do not substantially degrade car traffic conditions [66]. Under budget constraints, heuristic and metaheuristic methods are also widely used to design upgrade plans that target selected links while managing impacts on the wider network [67, 94].

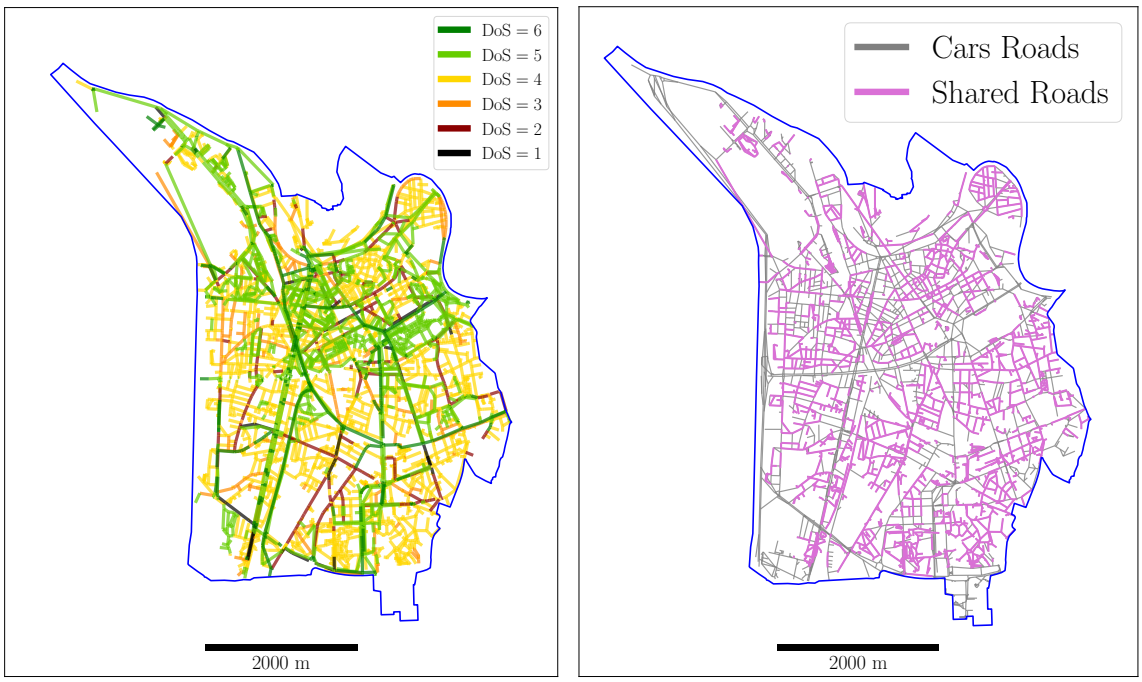

Figure 5.1: Multi-modal transportation network of Grenoble, France. *Left:* Cycling infrastructure classified by degree of separation (DoS) of bike lanes, ranging from 1 (least separated) to 6 (fully separated). *Right:* Road network for cars, including links shared with bicycles.

Other studies emphasize economic efficiency and broader societal impacts, for example by maximizing benefit–cost ratios or improving network connectivity under financial constraints [60, 61]. However, many existing formulations either (i) prioritize travel-time outcomes for bikes and cars without explicitly modeling cyclist safety and comfort, or (ii) focus on cycling-network improvements without systematically accounting for the traffic consequences of reallocating road capacity. In contrast, the framework proposed in this chapter explicitly integrates a safety-oriented objective for cyclists with traffic-performance considerations for automobiles within a unified bi-level model, enabling upgrade decisions that better reflect real-world multi-modal trade-offs.

5.2 Model Design

This section introduces the graph-based network model, car travel time modeling, and the bi-level optimization that jointly minimizes cyclist exposure and car travel time by steering bike flows toward safer routes.

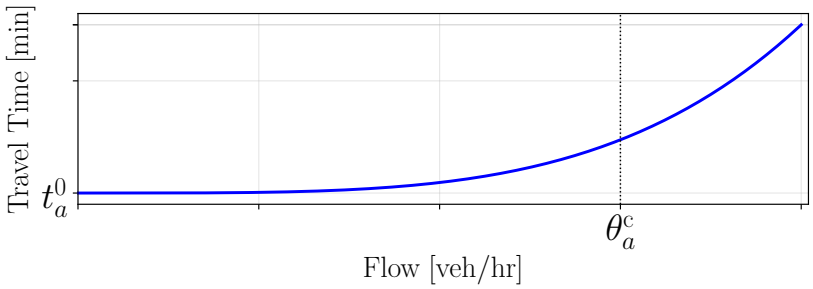

Figure 5.2: BPR function.

5.2.1 Graph Representation

As was presented in Chapter 2, now we consider the *multi-modal* transportation network. Unlike the previous chapters, where we focused solely on the cycling layer, here we explicitly model both bicycle and automobile movements in order to account for the impact of cycling infrastructure upgrades on car traffic. The transportation system is represented as a weighted directed graph (digraph)

$$\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{W}),$$

where $\mathcal{V}$ is the set of nodes (e.g., intersections, junctions, or road endpoints), $\mathcal{A}$ is the set of directed arcs representing road segments, and $\mathcal{W}$ is the set of associated arc weights. The network is composed of two distinct but potentially overlapping layers: the bicycle layer

$$\mathcal{G}^{\mathrm{b}} = \left(\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}}\right),$$

and the car layer

$$\mathcal{G}^{\mathrm{c}} = \left(\mathcal{V}^{\mathrm{c}}, \mathcal{A}^{\mathrm{c}}, \mathcal{W}^{\mathrm{c}}\right).$$

Each layer has its own node, arc, and weight sets, while the full network satisfies the union relations

$$\mathcal{V} = \mathcal{V}^{\mathrm{b}} \cup \mathcal{V}^{\mathrm{c}}, \qquad \mathcal{A} = \mathcal{A}^{\mathrm{b}} \cup \mathcal{A}^{\mathrm{c}}.$$

The layers may overlap, meaning that some intersections and road segments can be shared by both modes:

$$\mathcal{V}^{\mathrm{b}} \cap \mathcal{V}^{\mathrm{c}} \neq \emptyset, \qquad \mathcal{A}^{\mathrm{b}} \cap \mathcal{A}^{\mathrm{c}} \neq \emptyset.$$

For the car layer, each arc $a \in \mathcal{A}^{\mathrm{c}}$ is weighted by travel time, and we set $w_a^{\mathrm{c}} = t_a^{\mathrm{c}}$, since car links are represented through travel times and capacities rather than safety labels. Each car arc also has an associated geographical length, denoted d_a^{c}. When an arc is present in both layers (i.e., it represents the same physical road segment used by cyclists and cars), it belongs to $\mathcal{A}^{\mathrm{b}} \cap \mathcal{A}^{\mathrm{c}}$ and carries both a bicycle-related attribute (introduced in the previous chapters) and the car travel-time representation $(w_a^{\mathrm{c}}, d_a^{\mathrm{c}})$ defined here. To avoid heavy notation, we use the generic symbol a for an arc, and specify the layer only when needed through the superscript c. Finally, let $\mathcal{V}_{\mathcal{O}}^{\mathrm{c}}, \mathcal{V}_{\mathcal{D}}^{\mathrm{c}} \subseteq \mathcal{V}^{\mathrm{c}}$ denote the sets of origin and destination nodes for car trips.

5.2.2 Cars Travel Time

As in the bicycle network, let $q^{\mathrm{c}}_{(u,v)}$ denote the demand between nodes $(u,v) \in \mathcal{V}^{\mathrm{c}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{c}}_{\mathcal{D}}$. Unlike the all-or-nothing assignment, we adopt the UE traffic assignment principle [86], which assumes that no driver can reduce their travel time by unilaterally changing routes. Capturing this equilibrium requires accounting for capacity constraints: as demand increases or capacity decreases, travel times on a link increase, possibly making alternative paths more appealing. Let ϕ^{c}_a denote the car flow on arc $a \in \mathcal{A}^{\mathrm{c}}$, and $t^{\mathrm{c}}_a(\phi^{\mathrm{c}}_a)$ be the corresponding travel time. We use the Bureau of Public Roads (BPR) function [87]

$$t^{\mathrm{c}}_a(\phi^{\mathrm{c}}_a) = t^0_a \left(1 + \mu \left(\frac{\phi^{\mathrm{c}}_a}{\theta^{\mathrm{c}}_a}\right)^{\sigma}\right), \tag{5.1}$$

where t^0_a is the free-flow travel time, θ^{c}_a is the arc capacity, and $\mu = 0.15$, $\sigma = 4$ are fixed parameters. An illustration of this function is shown in Fig. 5.2. For low flow values, the BPR function remains relatively flat, and the travel time is close to the free-flow time t^0_a. As flow increases, the function increases non-linearly, indicating increased congestion and longer travel times. To compute equilibrium flows, let $\mathcal{P}^{\mathrm{c}}_{(u,v)}$ be the set of car paths between each $(u,v) \in \mathcal{V}^{\mathrm{c}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{c}}_{\mathcal{D}}$, and $\phi^{\mathrm{c}}_{p^{\mathrm{c}}_{(u,v)}}$ the path flow on $p^{\mathrm{c}} \in \mathcal{P}^{\mathrm{c}}_{(u,v)}$. The demand constraint reads:

$$\sum_{p^{\mathrm{c}} \in \mathcal{P}^{\mathrm{c}}_{(u,v)}} \phi^{\mathrm{c}}_{p^{\mathrm{c}}_{(u,v)}} = q^{\mathrm{c}}_{(u,v)}. \tag{5.2}$$

Total arc flow is obtained by summing over path flows:

$$\phi^{\mathrm{c}}_a = \sum_{(u,v)} \sum_{p^{\mathrm{c}} \in \mathcal{P}^{\mathrm{c}}_{(u,v)}} \phi^{\mathrm{c}}_{p^{\mathrm{c}}_{(u,v)}} \cdot \mathbb{I}_{a \in p^{\mathrm{c}}}, \tag{5.3}$$

where $\mathbb{I}_{a \in p^{\mathrm{c}}} = 1$ if arc a is part of path p^{c}, and 0 otherwise. Flows must be non-negative:

$$\phi^{\mathrm{c}}_{p^{\mathrm{c}}_{(u,v)}} \geq 0, \quad \forall p^{\mathrm{c}}, \forall (u,v), \tag{5.4}$$

$$\phi^{\mathrm{c}}_a \geq 0, \quad \forall a \in \mathcal{A}^{\mathrm{c}}. \tag{5.5}$$

Total car travel time is defined as the sum over all links of the product of flow and travel time:

$$\mathcal{TTT}^{\mathrm{c}}(\boldsymbol{\phi}^{\mathrm{c}}) = \sum_{a \in \mathcal{A}^{\mathrm{c}}} \phi^{\mathrm{c}}_a \cdot t^{\mathrm{c}}_a(\phi^{\mathrm{c}}_a), \tag{5.6}$$

where $\boldsymbol{\phi}^{\mathrm{c}} = [\phi^{\mathrm{c}}_a]^{\top}_{a \in \mathcal{A}^{\mathrm{c}}} \in \mathbb{R}^{|\mathcal{A}^{\mathrm{c}}|}_{\geq 0}$. To compute the user-equilibrium (UE) flow assignment, we use the Beckmann potential formulation [115]:

$$\mathcal{B}^{\mathrm{c}}(\boldsymbol{\phi}^{\mathrm{c}}) = \sum_{a \in \mathcal{A}^{\mathrm{c}}} \int_0^{\phi^{\mathrm{c}}_a} t^{\mathrm{c}}_a(x)\, \mathrm{d}x. \tag{5.7}$$

The UE traffic assignment can thus be formulated as:

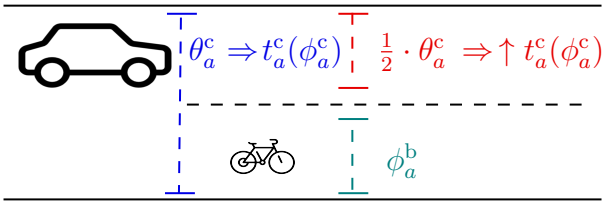

Figure 5.3: The creation/upgrading of a bike lane is modeled as reducing the capacity of the affected car link by half, resulting in an increased travel time $\uparrow t_a^c(\phi_a^c)$.

Problem 5.1 (UE Traffic Assignment).

$$\phi^{c\star} \in \mathbb{R}_{\geq 0}^{|\mathcal{A}^c|} = \arg \min_{\phi^c \in \mathbb{R}_{\geq 0}^{|\mathcal{A}^c|}} \mathcal{B}^c(\boldsymbol{\phi}^c)$$

s.t. Eq. (5.2) – (5.5).

Then, $\boldsymbol{\phi}^{c\star} = \left[\phi_a^{c\star}\right]_{a\in\mathcal{A}^c}^\top$ represents the UE flow assignment, in which no driver can unilaterally reduce their travel time. This flow pattern satisfies the conditions of a Nash equilibrium, as defined in the classical traffic assignment model. The convexity of the Beckmann potential $\mathcal{B}^c$ and its connection to user equilibrium are formally established in [116]. In the next section, we integrate this car traffic assignment model with the cyclist exposure model to form the full bi-level optimization framework.

The following theorem uses the convexity of the potential function to establish the correspondence between its minimizer and Nash equilibrium.

Theorem 5.1: (Beckmann Potential and User Equilibrium).

The Beckmann potential function:

$$\mathcal{B}^c(\boldsymbol{\phi}^c) = \sum_{a\in\mathcal{A}^c} \int_0^{\phi_a^c} t_a^c(x)\,\mathrm{d}x \tag{5.8}$$

- *i)* is convex,
- *ii)* and every minimizer $\boldsymbol{\phi}^{c\star}$ of Problem 5.1 is a user equilibrium, i.e., a Nash equilibrium.

5.2.3 Interaction Between Cars and New Bike Lanes

Upgrading bike lanes can reduce the capacity of adjacent car lanes. However, as shown in [117], capacity effects are assumed negligible unless the upgraded facility is fully separated

(DoS 5–6). In our setting, only fully separated bike lanes, i.e., those with $\tilde{w}_a^{\mathrm{b}} \geq 5$, affect car capacity. The capacity reduction is modeled as:

$$\tilde{\theta}_a^{\mathrm{c}} = \begin{cases} \frac{1}{2}\,\theta_a^{\mathrm{c}}, & \text{if } a \in \mathcal{A}^{\mathrm{c}} \cap \mathcal{A}^{\mathrm{b}} \text{ and } \tilde{w}_a^{\mathrm{b}} \geq 5, \\ \theta_a^{\mathrm{c}}, & \text{otherwise.} \end{cases} \tag{5.9}$$

where $\tilde{\theta}_a^{\mathrm{c}}$ is the updated car capacity on arc a, θ_a^{c} is the original capacity, and $\tilde{w}_a^{\mathrm{b}}$ is the DoS of the upgraded bike lane on the shared arc. Following [62], the $\frac{1}{2}$ factor reflects that a standard bike lane occupies roughly half the width of a car lane, thus reducing vehicular capacity by about 50% on the affected lane. To keep the bi-level design problem computationally tractable in this first multi-modal formulation, we adopt a fixed-flow (pre-upgrade UE) approximation: the car flow vector $\boldsymbol{\phi}^{\mathrm{c}^\star}$ is computed once from Problem 5.1 using the original capacities and is then held fixed during the optimization. Under this assumption, drivers do not re-route in response to capacity changes induced by cycling upgrades; only link travel times $t_a^{\mathrm{c}}(\cdot)$ increase through the updated capacities $\tilde{\theta}_a^{\mathrm{c}}$ in Eq. (5.1). This is a strong simplification, since in reality some drivers may shift routes to avoid newly congested links, and such re-routing can create non-local effects on travel times. We therefore interpret this model as a conservative, first-step screening tool: it captures the direct travel-time penalty on the car network due to capacity loss, but it does not capture network-wide adaptation. To assess the impact of this assumption, we perform a post-optimization validation in which we recompute the Wardrop user equilibrium on the upgraded network and compare pre- and post-upgrade assignments. Moreover, Chapter 6 relaxes the fixed-flow approximation within the optimization loop by introducing a graph neural-network surrogate trained on UE solutions. This surrogate provides fast predictions of near-equilibrium car flows and travel times under changing capacities (and demand), enabling us to approximate UE re-routing effects without solving a full equilibrium problem at every iteration.

5.2.4 Optimization Problem

Given a fixed budget B, the objective is to upgrade the cycling infrastructure by increasing the DoS on selected segments, minimizing cyclist exposure while limiting car travel time increases. Car flows $\phi^{\mathrm{c}^\star}$ remain fixed during optimization, obtained with the pre-upgrade UE assignment from Problem 5.1. However, travel time penalties from high-DoS upgrades are accounted for via updated capacities following Eq. (5.9). Each upgrade incurs a cost and must comply with a global budget constraint.

- $\boldsymbol{w}^{\mathrm{b}} = \left[w_a^{\mathrm{b}}\right]_{a\in\mathcal{A}^{\mathrm{b}}}^{\top} \in \mathbb{N}^{|\mathcal{A}^{\mathrm{b}}|}$: vector of original DoS values for each arc in the bike network;
- $\tilde{\boldsymbol{w}}^{\mathrm{b}} = \left[\tilde{w}_a^{\mathrm{b}}\right]_{a\in\mathcal{A}^{\mathrm{b}}}^{\top} \in \Omega$: vector of decision variables (candidate configuration) representing the upgraded DoS values, where Ω is the constraint set that defined as $\Omega = \left\{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \mathbb{N}^{|\mathcal{A}^{\mathrm{b}}|} \colon w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6, \quad \forall a \in \mathcal{A}^{\mathrm{b}}\right\}$;

Additionally, we define the following variables, all normalized between their respective minimum and maximum values: Now, we present the optimization problem that consists of two

Symbol	Description	Domain	Unit
$\phi_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Flow on arc $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	bike/hr
d_a^{b}	Length of $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	m
$\bar{d}_a^{\mathrm{b}}$	Cyclist distance of $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	m
$\mathcal{E}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Cyclist exposure on $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	-
$\mathcal{TTT}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Total cars travel time	$\mathbb{R}_{\geq 0}$	hr
$\tilde{\theta}_a^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Updated car lane capacity on $a \in \mathcal{A}^{\mathrm{c}}$	$\mathbb{R}_{\geq 0}$	veh/hr
$\phi_a^{\mathrm{c}\star}$	Pre-upgrade car flow on lane $a \in \mathcal{A}^{\mathrm{c}}$ obtained by solving Problem 5.1	$\mathbb{R}_{\geq 0}$	veh/hr
$\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$	Upgrade cost for $a \in \mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	€

Table 5.1: Abbreviated notation for variables in cycling network model, writing $(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ next to a variable indicates that it is a function of the candidate configuration for bike lane improvements.

interconnected levels:

Problem 5.2 (Trade-Off Network-Upgrade).

$\mathbb{U}_{\text{pper}}$:

$$\tilde{\boldsymbol{w}}^{\mathrm{b}\star} = \arg\min_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega} \; \gamma \cdot \left(\underbrace{\sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot f_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot d_a^{\mathrm{b}}}_{\text{Cyclist exposure } \mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}})} \right)$$

$$+ (1-\gamma) \cdot \left(\underbrace{\sum_{a \in \mathcal{A}^{\mathrm{c}}} \phi_a^{\mathrm{c}\star} \cdot t_a^0 \cdot \left(1 + \mu \left(\frac{\phi_a^{\mathrm{c}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{\tilde{\theta}_a^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})} \right)^{\sigma} \right)}_{\text{Total car travel time}} \right) \tag{5.10}$$

$$\text{s.t.} \sum_{a \in \mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \leq B, \tag{5.11}$$

$$\tilde{\theta}_a^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \begin{cases} \frac{1}{2}\,\theta_a^{\mathrm{c}}, & \text{if } a \in \mathcal{A}^{\mathrm{c}} \cap \mathcal{A}^{\mathrm{b}} \text{ and } \tilde{w}_a^{\mathrm{b}} \geq 5, \\ \theta_a^{\mathrm{c}}, & \text{otherwise.} \end{cases}; \tag{5.12}$$

$\mathbb{L}_{\text{ower}}$:

$$\boldsymbol{\phi}^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \arg\min_{\phi^{\mathrm{b}} \in \mathbb{R}_{\geq 0}^{|\mathcal{A}^{\mathrm{b}}|}} \sum_{a \in \mathcal{A}^{\mathrm{b}}} \bar{d}_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \phi_a^{\mathrm{b}}, \tag{5.13}$$

$$\text{s.t. } \phi_a^{\mathrm{b}} = \sum_{(u,v) \in \mathcal{V}_{\mathcal{O}}^{\mathrm{b}} \times \mathcal{V}_{\mathcal{D}}^{\mathrm{b}}} q_{(u,v)}^{\mathrm{b}} \cdot \mathbb{I}_{a \in p_{(u,v)}^{\mathrm{b}}} \geq 0, \; \forall a \in \mathcal{A}^{\mathrm{b}}. \tag{5.14}$$

The *Upper-Level Problem* ($\mathbb{U}_{\text{pper}}$) optimizes the distribution of infrastructure upgrades. The first term in Eq. (5.10) reduces cyclist exposure by prioritizing links with high flow, traffic, and length. The second term limits additional car travel time due to capacity losses on shared links with DoS $\geq$ 5. The trade-off is managed by the weight $\gamma \in [0, 1]$. Constraints Eq. (5.11) and Eq. (5.12) enforce the budget limit and car lane capacity reduction rule.

The *Lower-Level Problem* ($\mathbb{L}_{\text{ower}}$) computes bike flows assuming cyclists use shortest paths based on cyclist distance. The resulting flows $\boldsymbol{\phi}^{\mathrm{b}^\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ influence the upper-level objective via exposure, thus coupling both levels. Solving Problem 5.2 is computationally challenging. The upper level entails a discrete selection of DoS upgrades under a global budget constraint, which already contains the structure of a 0–1 knapsack problem as a special case. This observation implies NP-hardness, and closely parallels the hardness arguments used for related graph-upgrade formulations (see, e.g., [53]). In addition, the objective couples these discrete upgrade decisions with network-wide flow responses through the lower-level assignment and with nonlinear congestion penalties on the car layer via the BPR travel-time function. As a result, the overall problem combines combinatorial choice with nonlinear, network-dependent effects, which makes exact solution methods impractical at city scale and motivates the use of heuristic approaches.

5.3 Case Study

In this section, we present a case study based in Grenoble, France, and analyze the optimization results obtained using the GA-U and MLCC algorithms to solve Problem 5.2. Additionally, we assess the impact of bike lane upgrades on car travel times using a post-upgrade UE traffic assignment.

5.3.1 Problem Parameters and Optimal Solution

To solve the all-or-nothing traffic assignment problem and determine cyclist flows, as described in Chapter 4.2.1, and similarly to the previous chapters, let $B = \alpha \cdot B_{\max}$ be the available budget, where $B_{\max}$ is the cost to upgrade the full network to DoS 6 and $\alpha \in [0, 1]$. Upgrade costs per meter depend on DoS levels, based on French public estimates [110] as in Eq. (3.43). Fig. 5.4 shows the normalized car flows $\phi^{\mathrm{c}^\star}$ obtained via the Frank–Wolfe algorithm [118, 119] for Problem 5.1, which serves as the baseline (pre-upgrade) traffic assignment. These flows are then used to solve Problem 5.2 across three budget levels $\alpha \in \Psi = \left\{\frac{1}{3}, \frac{2}{5}, \frac{1}{2}\right\}$ and eleven values of the trade-off parameter $\gamma \in \Lambda = \{0, \frac{1}{10}, \ldots, 1\}$. For each α, we obtain a discrete set of Pareto points

$$\left\{\left(\gamma \cdot \mathcal{E},\ (1-\gamma) \cdot \mathcal{TTT}^{\mathrm{c}}\right)\right\}_{\gamma \in \Lambda}, \tag{5.15}$$

which characterizes the trade-off between cyclist exposure $\mathcal{E}$ and the total travel time of cars $\mathcal{TTT}^{\mathrm{c}}$. As expected, $\gamma = 0$ prioritizes car travel time, whereas $\gamma = 1$ prioritizes cyclist exposure. Choosing a single "best" point on this trade-off set is inherently a policy decision and depends on stakeholder preferences, risk tolerance, and local priorities; hence it should

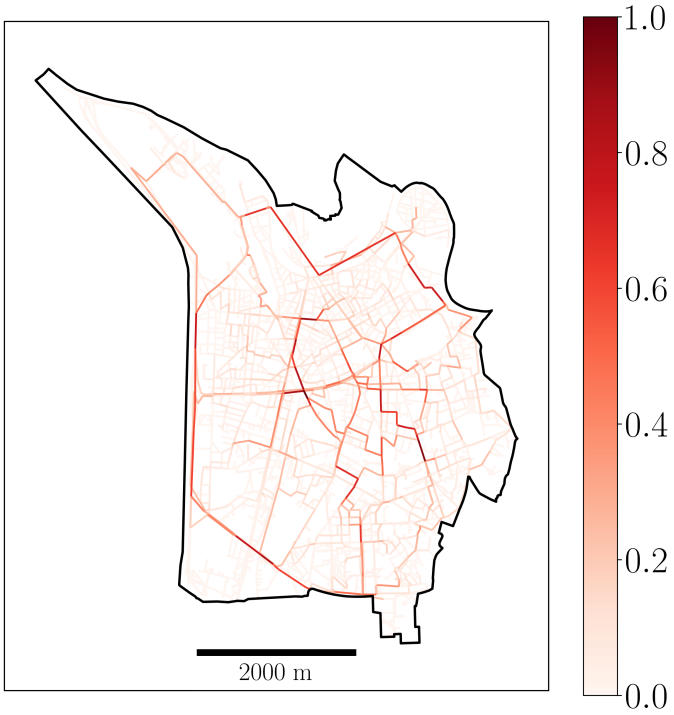

Figure 5.4: Original pre-upgrade UE flows (normalized) obtained by solving Problem 5.1.

ultimately be determined by a decision maker rather than by the model alone. Nevertheless, for the purpose of producing a concise quantitative comparison across budgets in this academic study, we adopt a simple rule and select the point closest to the origin in the normalized objective space:

$$\gamma_\alpha^\star = \arg\min_{\gamma \in \Lambda} \sqrt{(\gamma \cdot \mathcal{E})^2 + ((1-\gamma) \cdot \mathcal{TTT}^{\mathrm{c}})^2}. \tag{5.16}$$

Fig. 5.5 reports the resulting trade-offs. Higher budgets (e.g., $\alpha = \frac{1}{2}$) generally enable lower cyclist exposure at the cost of a modest increase in car travel time. Fig. 5.5 presents the normalized outcomes for each γ, illustrating the expected pattern: as γ increases, exposure decreases while travel time increases. The selected reference points $\gamma_\alpha^\star$ are highlighted in red. The figure also provides information on the relative convergence speed associated with each reference point. Among them, $\gamma_{\frac{1}{2}}^\star$ appears to reach a stable high-quality solution more rapidly. One possible explanation is that, for larger budget levels, the optimization is less constrained by feasibility and the algorithm has access to a broader set of admissible upgrade combinations that yield comparable objective improvements. This makes it easier for the population to identify and propagate promising building blocks early in the search. By contrast, under tighter budget constraints, the number of feasible configurations that simultaneously satisfy the budget and produce meaningful improvement is smaller, so the search becomes more selective and requires more generations to identify competitive solutions. Hence, the slower convergence observed for lower budget levels may reflect not weaker search performance, but a more constrained combinatorial structure of the decision space.

5.3.2 Impact on Car Routing Choice and Travel Time

To assess the impact of upgrades at $\gamma_\alpha^\star$, we compare the UE traffic assignment computed on the pre-upgrade network (i.e., assuming drivers maintain their original routes) with the post-upgrade UE assignment, where drivers re-optimize their routes in response to updated capacities. Using the Frank–Wolfe method, Fig. 5.7 shows that, after re-assigning routes on the post-upgrade road network, car travel times increase only moderately. Under the GA-U solutions, the maximum increase is 11.1% (for $\alpha = \frac{1}{2}$), suggesting that the proposed upgrades improve cycling conditions while causing limited disruption to automobile traffic.

Why do the maps look similar across α? Although Fig. 5.8 (a)–(c) may appear visually close at the city scale, the upgrade plans are materially different across budgets: smaller budgets concentrate upgrades on a few high-impact segments (typically along the most-used corridors), while larger budgets allow upgrades to be spread across additional links and alternative routes. The differences are therefore most visible in which corridors are upgraded and how far upgrades extend, rather than in a global change of the full map color distribution. Moreover, because the objective is exposure-based and flow-weighted, large gains can often be obtained by upgrading a relatively small set of bottleneck links early on; once these are improved, additional budget yields more incremental improvements, which can make the visual contrast between $\alpha = \frac{1}{3}$, $\frac{2}{5}$ and $\frac{1}{2}$ appear modest.

Normalization and the exposure–travel-time trade-off. The bi-objective formulation uses an affine combination between cyclist exposure and car total travel time. Since these quantities have different units and scales, both components are normalized prior to aggregation (so that neither term dominates purely due to magnitude), and γ controls the relative importance of the normalized components. Consequently, the limited variation across α should not be interpreted as an "over-weight" of exposure, but rather as (i) the saturation effect described above, and (ii) the fact that the GA-U solutions are constrained to avoid large capacity losses on the car network.

In contrast, the MLCC-based plans tend to affect car traffic more. This behavior is expected from the design of MLCC: it explicitly seeks to maximize the size of the largest connected component of DoS = 6 infrastructure. Reaching DoS 6 typically requires fully segregated facilities, which in practice often implies reallocating roadway space (e.g., converting a motor-vehicle lane or removing on-street parking) along continuous corridors. Such corridor-scale interventions reduce car capacity over longer stretches rather than at isolated links, thereby creating larger and more systematic impacts on car travel times once traffic is reassigned.

This contrast is visible in Fig. 5.7: MLCC produces substantially higher $\mathcal{TTT}^{\mathrm{c}}$ than GA-U for the same α, especially under the fixed-route (pre-upgrade) car assignment, because MLCC does not internalize the objective of limiting car impacts. The post-upgrade UE reassignment (red bars) reduces the discrepancy by allowing drivers to re-route, but MLCC

remains consistently worse in car travel time. The flow maps in Fig. 5.9 (d)–(f) also show larger and more widespread changes in car-flow patterns than the GA-U solutions, consistent with longer contiguous capacity reductions. Overall, the main qualitative difference is therefore not between nearby α values, but between the two upgrade logics: GA-U explicitly balances cycling safety improvements with car-network impacts, whereas MLCC prioritizes DoS 6 connectivity and can induce substantially larger car-delay costs.

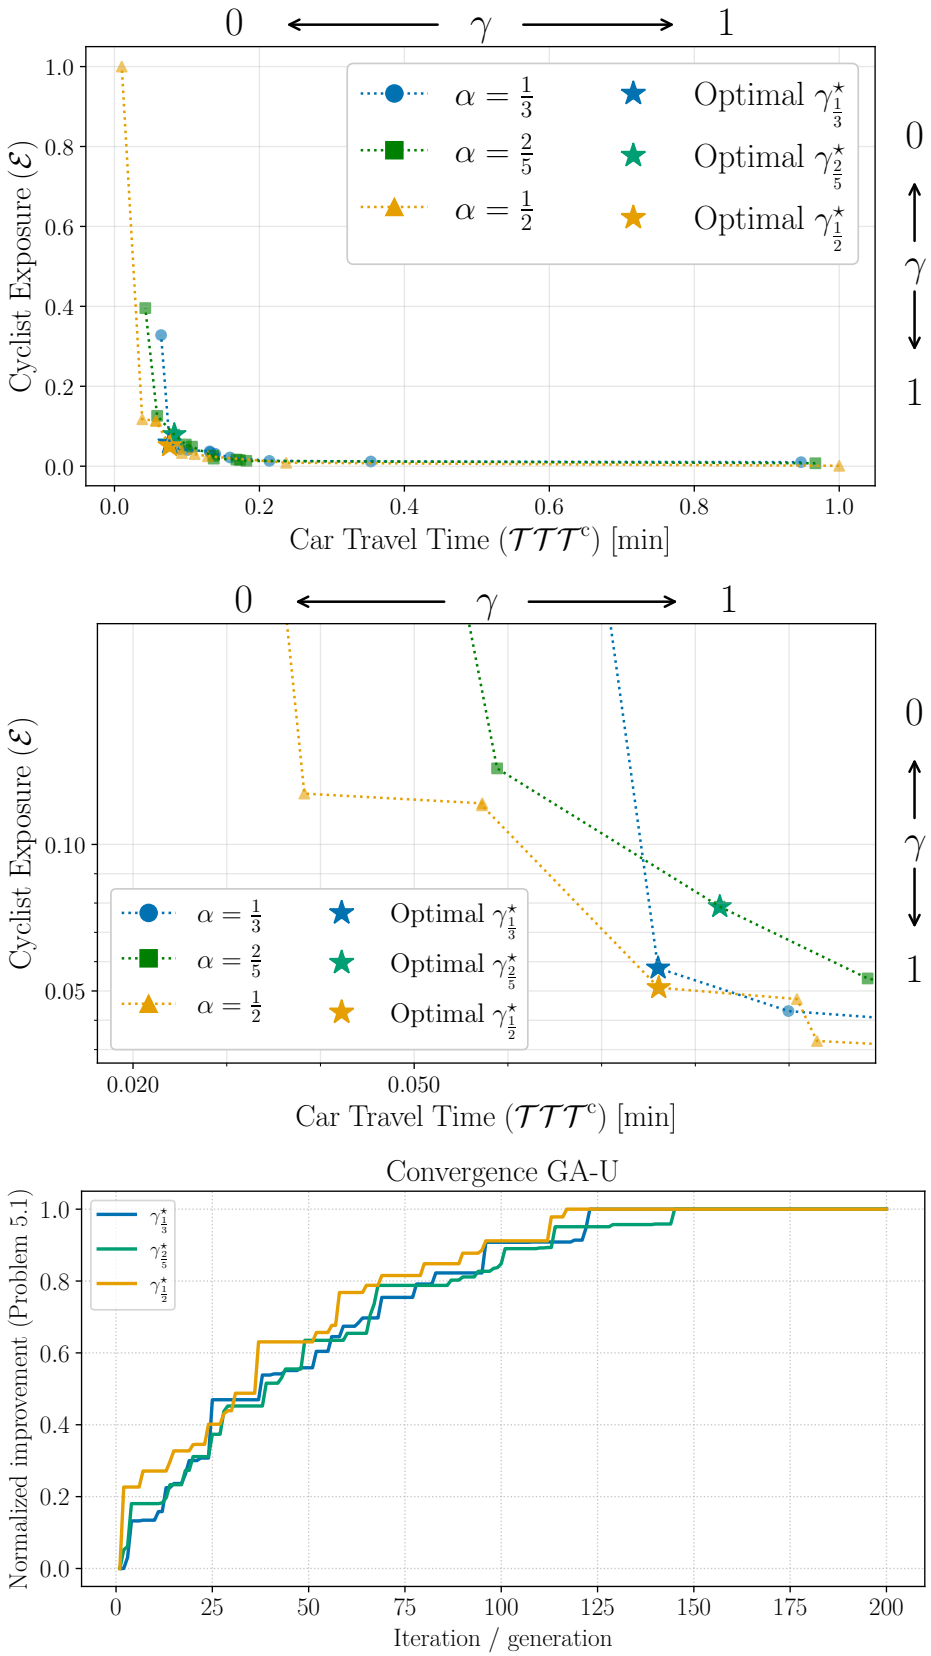

Figure 5.5: Normalized travel time versus cyclist exposure. The top plot displays the Pareto fronts resulting from the trade-off Problem 5.2 across different values of $\gamma \in \Lambda$ and $\alpha \in \Psi$. The medium plot presents a zoomed-in view focusing on the optimal γ-constraint solutions. The bottom figure present the convergence analysis of the GA-U algorithm for solving Problem 5.2 on the optimal points $\gamma^\star_\alpha$.

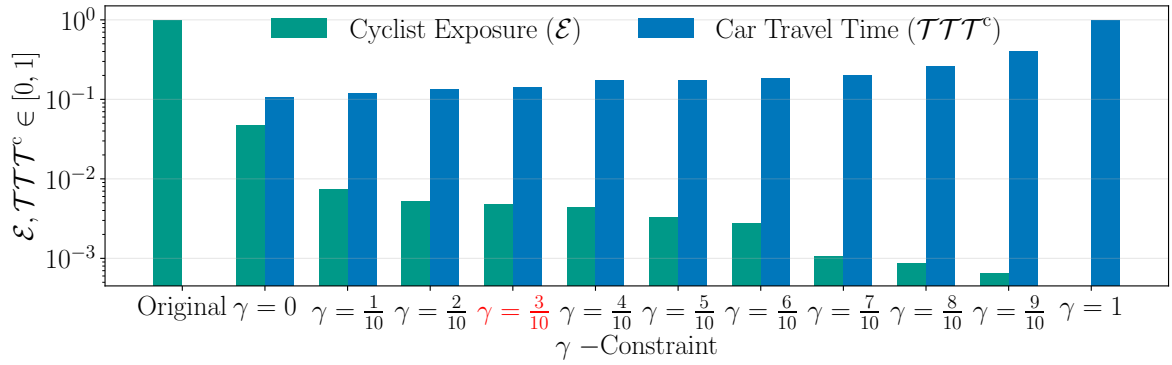

(a) $\alpha = \frac{1}{3}$

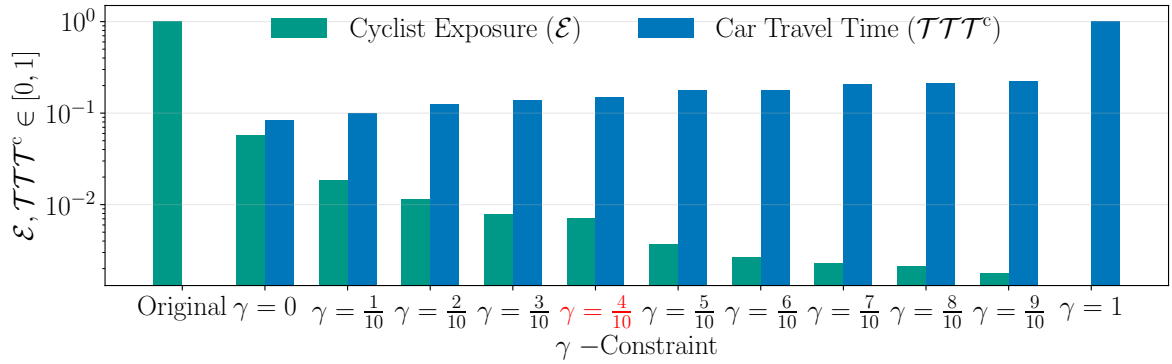

(b) $\alpha = \frac{2}{5}$

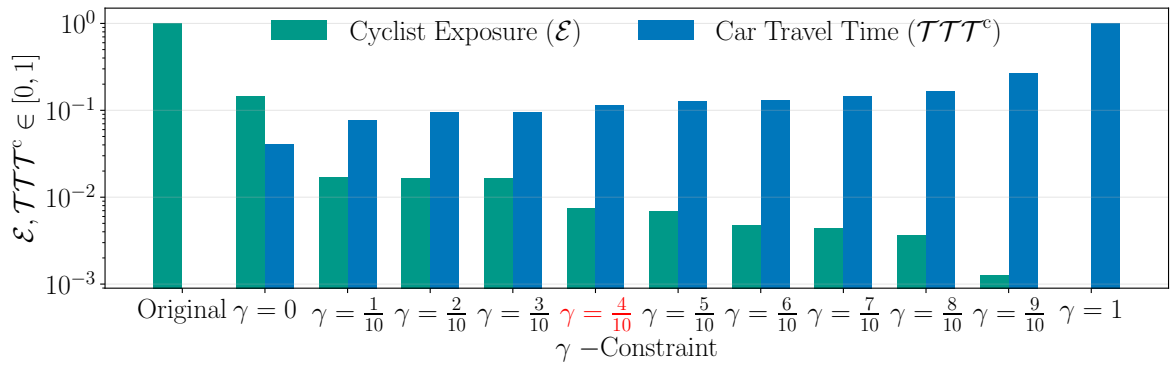

(c) $\alpha = \frac{1}{2}$

Figure 5.6: Normalized cyclist exposure and travel time for each $\gamma \in \Lambda$ and $\alpha \in \Psi$. In red color are the $\gamma^\star_{\frac{1}{3}}, \gamma^\star_{\frac{2}{5}}$ and $\gamma^\star_{\frac{1}{2}}$ from top to bottom, respectively.

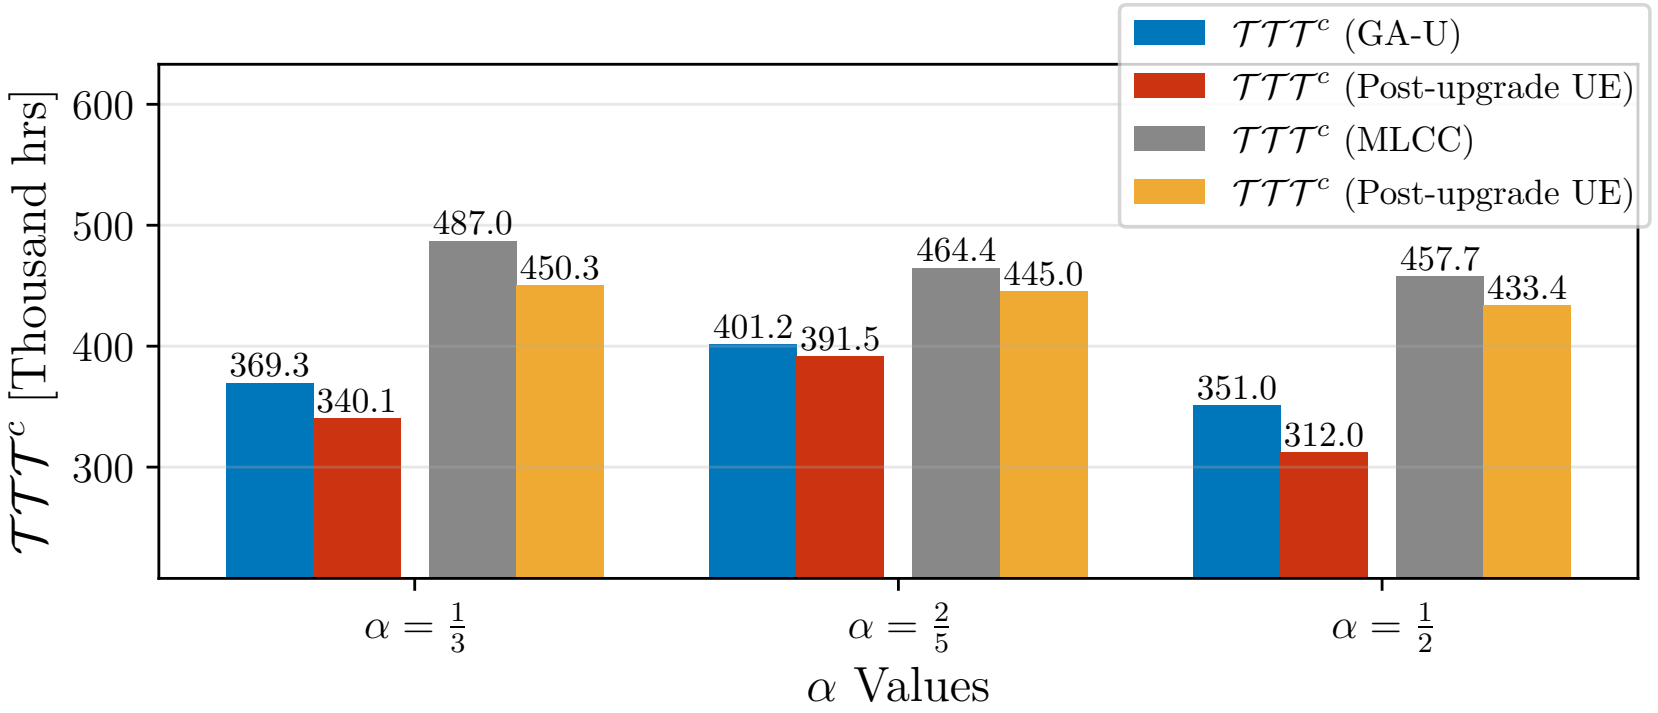

Figure 5.7: Comparison of car travel times at the optimal trade-off points $\gamma_\alpha^\star$, with post-upgrade UE traffic assignment after bike lane upgrades at these locations.

(a) Optimal upgrade for $\gamma^{\star}_{\frac{1}{3}}$

(b) Optimal upgrade for $\gamma^{\star}_{\frac{2}{5}}$

(c) Optimal upgrade for $\gamma^{\star}_{\frac{1}{2}}$

(d) Post-upgrade user equilibrium traffic assignment for $\gamma^{\star}_{\frac{1}{3}}$

(e) Post-upgrade user equilibrium traffic assignment for $\gamma^{\star}_{\frac{2}{5}}$

(f) Post-upgrade user equilibrium traffic assignment for $\gamma^{\star}_{\frac{1}{2}}$

Figure 5.8: GA-U upgrades and post-upgrade user-equilibrium (UE) assignments for three budgets.

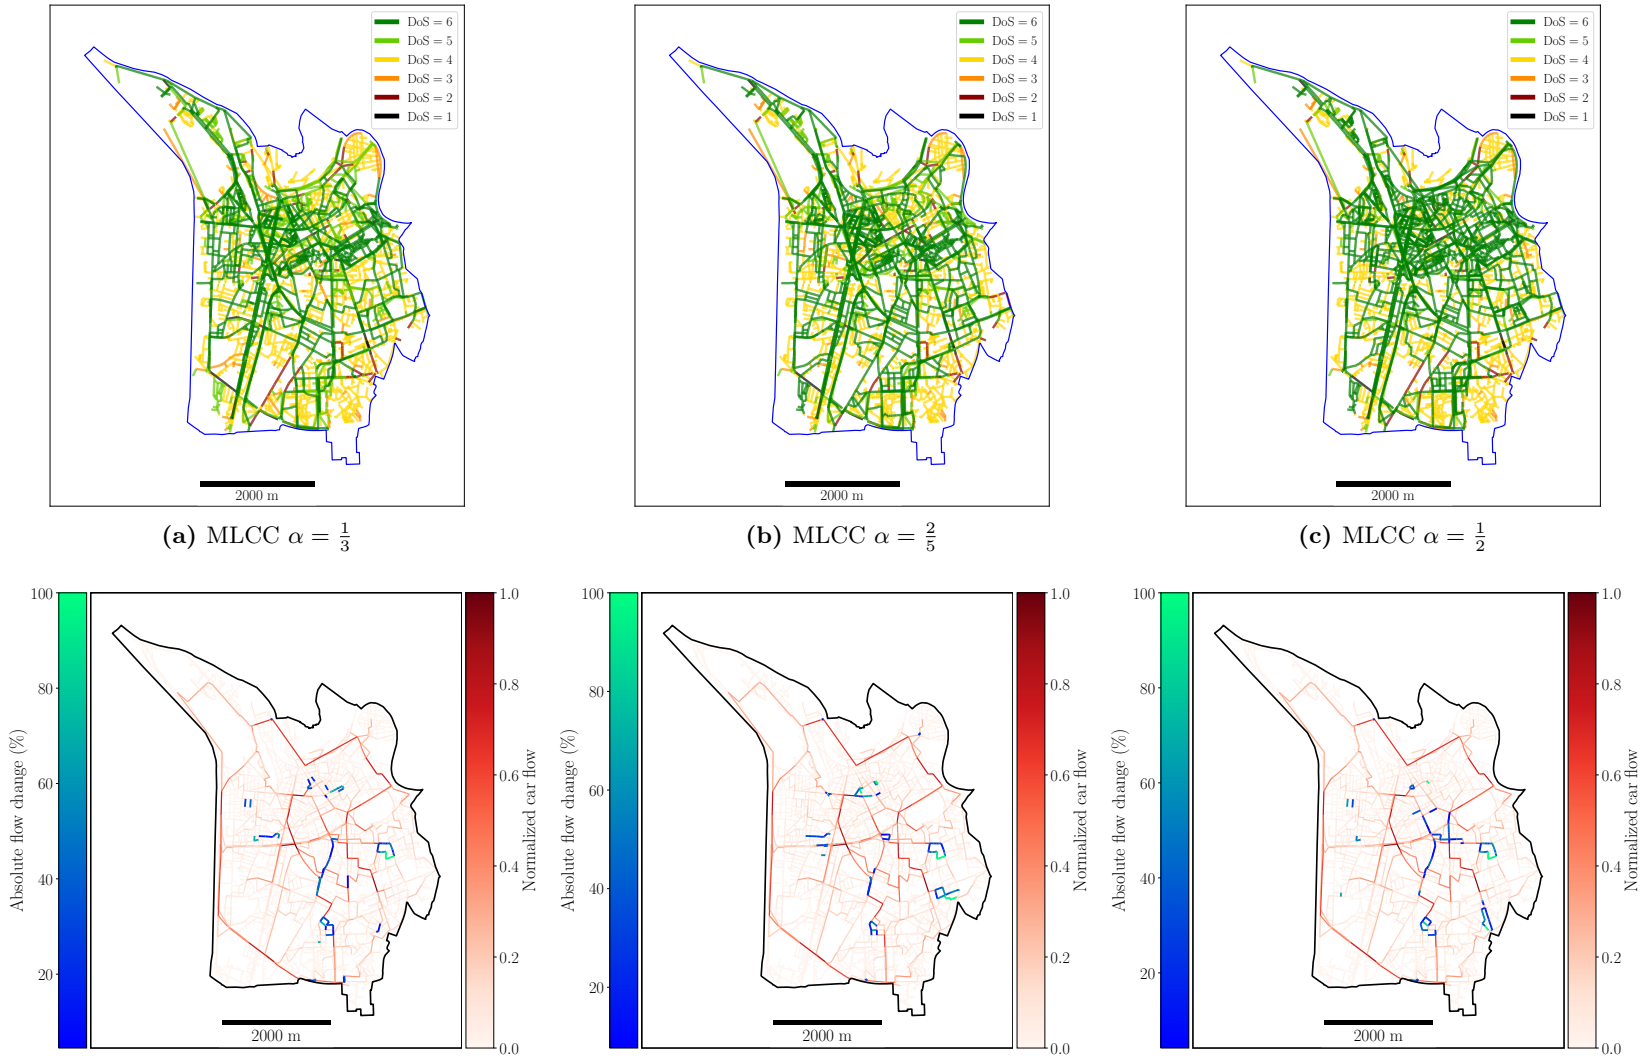

(a) MLCC $\alpha = \frac{1}{3}$

(b) MLCC $\alpha = \frac{2}{5}$

(c) MLCC $\alpha = \frac{1}{2}$

(d) Post-MLCC upgrade user equilibrium traffic assignment for $\alpha = \frac{1}{3}$

(e) Post-upgrade user equilibrium traffic assignment for $\alpha = \frac{2}{5}$

(f) Post-MLCC upgrade user equilibrium traffic assignment for $\alpha = \frac{1}{2}$

Figure 5.9: MLCC upgrades and post-upgrade user-equilibrium (UE) assignments for three budgets.

5.4 Concluding Remarks

This chapter presents a bi-objective optimization framework for cycling-infrastructure planning that explicitly captures a central urban trade-off: improving cycling safety while limiting the deterioration of car-network performance. Using Grenoble as a case study, we model cycling safety through a flow-weighted exposure indicator and represent motor-traffic impacts through total car travel time under a capacity-sensitive user-equilibrium (UE) assignment. We solve the resulting design problem with the GA-U heuristic and, across multiple budget levels, obtain Pareto frontiers that quantify how exposure reductions translate into (typically modest) car travel-time increases. Importantly, these conclusions remain consistent when car routes are recomputed after the upgrade using a post-upgrade UE assignment, indicating that the identified plans are not artifacts of fixed-route assumptions but remain robust under realistic traffic re-routing.

Beyond the numerical results, the main message is methodological: multi-modal planning decisions should be evaluated as system interactions rather than mode-specific improvements in isolation. The proposed framework provides planners with interpretable trade-off curves showing what level of safety improvement can be achieved for a given tolerance to car-delay impacts—and highlights that meaningful exposure reductions can often be obtained without severe disruption to motor traffic, particularly when investments target links that improve safety along high-demand cycling corridors while avoiding large capacity losses on critical car bottlenecks.

At the same time, this chapter intentionally adopts a "first-step" representation of demand: cycling demand is fixed and the car demand is exogenous, so the framework quantifies network-performance trade-offs conditional on current travel patterns. In practice, however, safer cycling conditions can induce modal shifts from cars to bicycles, and the acceptability of travel-time impacts is not purely an efficiency question but also a distributional one: who experiences increased travel times, and whether minimum accessibility standards remain satisfied across neighborhoods and user groups.

For these reasons, the next chapter extends the framework in two directions. First, we incorporate a transport mode choice model to make demand endogenous and capture the feedback loop between infrastructure improvements and travelers' decisions to switch from car to bicycle. Second, we move from aggregate utilitarian perspective to a transport justice-oriented perspective by formulating both utilitarian and sufficientarian problems, where the objective is not only to improve average safety and efficiency but also to ensure that minimum thresholds of safety (for cyclists) and accessibility (for drivers and cyclists) are met. Together, these extensions transform the initial multi-modal trade-off analysis into a richer decision-support tool that reflects both behavioral adaptation and normative planning goals.

Chapter 6

Network Design with Flow Redistribution in a Multi-Modal Network for Transport Justice

Contents

6.1 Introduction . . . 103
6.1.1 Summary of Results . . . 104
6.1.2 Literature Review . . . 105
6.2 Model Design . . . 107
6.2.1 Total system travel time . . . 107
6.2.2 Interaction Between Cars and New Bike Lanes . . . 107
6.2.2.1 Mode Choice Model Based on Temporal Data . . . 108
6.2.2.2 Surrogate-Based Flow Prediction with Feasibility Projection . . . 111
6.2.3 Optimization Problem (Utilitarian Perspective) . . . 112
6.2.4 Optimization Problem (Sufficientarian Perspective) . . . 114
6.3 Case Study . . . 118
6.3.1 Algorithm Comparisons . . . 118
6.4 Concluding Remarks . . . 128

6.1 Introduction

In the previous chapter, we introduced a multi-modal framework to model the interaction between cycling-infrastructure upgrades and changes in car travel times resulting from capacity reallocation. In this chapter, we extend that setting in two directions that are essential for policy relevance. First, we incorporate *modal shift*: improvements to cycling safety and comfort can induce a fraction of drivers to switch from car to bicycle, which in turn changes demand levels, redistributes flows, and may partially offset congestion effects on the road network. Second, we explicitly address *transport justice* through a sufficientarian perspective, recognizing that infrastructure programs should not only improve average conditions, but also ensure that minimum standards of safety and accessibility are met for all users.

Physically separated cycling facilities are known to enhance perceived safety and increase ridership [4, 44]. At the same time, reallocating road space to bicycles may reduce vehicle capacity and increase car travel times. The potential for modal shift complicates this trade-off: while some upgrades may initially worsen car performance, they can also encourage a transfer of trips to cycling, thereby reducing car demand and mitigating travel-time impacts. Capturing these feedback effects requires a multi-modal model that jointly accounts for (i) cyclist route choice and exposure, (ii) car traffic assignment, and (iii) demand shifts between modes.

Beyond efficiency, cycling investments raise distributional questions: who benefits from safety improvements, and who bears the costs of longer travel times? Building on transport justice theory [7], we consider two complementary perspectives. A *utilitarian* perspective seeks to improve overall system performance, for instance by reducing aggregate cyclist exposure and total travel time. A *sufficientarian* perspective, in contrast, focuses on ensuring that all users achieve at least a minimum standard, e.g., that cyclists can access routes meeting a safety threshold and that drivers are not pushed beyond an unacceptable travel-time burden. Incorporating these principles into network design motivates optimization formulations that combine system-wide objectives with explicit safeguards for vulnerable or disadvantaged users.

This chapter presents a multi-modal planning framework that integrates induced modal shift and justice considerations under budget constraints. The model updates travel demand across modes as cycling infrastructure is upgraded, and evaluates traffic responses through a computationally efficient surrogate model that approximates user-equilibrium car flows at city scale. We then introduce two bi-level optimization formulations: (i) a utilitarian model that balances cyclist exposure and car travel time, and (ii) a sufficientarian model that prioritizes meeting minimum safety and travel-time standards, while still improving overall performance.

6.1.1 Summary of Results

Applied to the Grenoble network, the multi-modal bi-level framework consistently finds upgrade plans that improve cycling safety while keeping impacts on automobile performance within an acceptable range once traffic re-equilibration is taken into account. Across the tested budget levels, the selected interventions reduce cyclist exposure and strengthen access to low-stress cycling routes, showing that safety gains are not confined to isolated links but can be achieved in a way that improves the usability of the network as a whole.

A central result is that introducing modal shift changes the assessment of "disruption" to car traffic. Some upgrades that would mechanically reduce road capacity and increase car travel times become more feasible once demand is allowed to respond: as cycling conditions improve, a fraction of trips switches from car to bicycle, lowering car demand and partially offsetting congestion effects through the new user-equilibrium flow pattern. In Grenoble, this feedback makes it possible to obtain substantial safety improvements with smaller-than-expected travel-time penalties compared with a fixed-demand setting.

The comparison between the utilitarian and sufficientarian formulations then highlights

the difference between improving averages and improving distributions. The utilitarian model performs well in terms of aggregate efficiency: it identifies upgrade portfolios that achieve large exposure reductions per unit of budget while limiting the overall increase in car travel time. Because it optimizes system-level totals, however, it can still leave some users below acceptable standards, e.g., cyclists whose feasible routes remain relatively unsafe, or drivers concentrated on corridors where travel-time impacts are more pronounced.

By construction, the sufficientarian formulation produces the most balanced outcomes. In the Grenoble case study it reduces the share of cyclists who remain "safety-insufficient" by expanding the availability of routes that satisfy the safety threshold, and at the same time it limits the fraction of drivers pushed beyond a maximum acceptable travel-time burden. This leads to solutions that distribute benefits more evenly across the city, prioritizing interventions that close gaps in safety and accessibility rather than only maximizing the largest aggregate gains. Importantly, these fairness improvements are achieved without abandoning efficiency: the sufficientarian strategy still delivers meaningful exposure reductions and maintains competitive overall performance, but with clearer guarantees on minimum standards.

Overall, the results show that cycling-upgrade planning can be made both operationally realistic and policy-relevant by jointly modeling cyclist routing, car traffic equilibrium, and induced demand shifts, and by embedding explicit justice criteria. The surrogate UE model enables city-scale optimization, and the Grenoble experiments illustrate that fairness-oriented design can improve safety and accessibility while keeping efficiency losses limited supporting the case that "just" low-carbon planning does not require sacrificing system performance.

6.1.2 Literature Review

Our work draws on three primary domains: cycling network optimization, the analysis of modal shifts in urban mobility and transport justice.

Cycling Network Optimization: Research on cycling network planning addresses objectives such as usage, connectivity, safety, and multi-modal impacts. Early studies optimized network design to increase bike-lane use—measured by the total distance traveled on designated facilities—as a proxy for infrastructure utilization and accessibility, while improving topological connectivity [56, 57, 60]. Subsequent works examined benefit–cost trade-offs and proposed budget-constrained heuristics to guide investment decisions [61, 67, 94]. More recent graph-theoretic approaches explicitly targeted safety by reconfiguring networks toward higher protection levels, though they often overlook how improvements are distributed across users and regions [54]. In parallel, multi-modal interactions have been modeled through bi-level formulations that balance cycling upgrades with car accessibility [56, 57, 62] and heuristic methods that capture traffic disruption under lane reallocation [66]. Overall, most approaches emphasize efficiency or travel time. Safety is sometimes considered at an aggregate level, but the fairness of how safety and accessibility gains are distributed, who benefits, where, and by how much, remains largely underexplored.

Modal Shift in Urban Mobility: Utility cycling has become an increasingly credible alter-

native to driving in urban areas, although cars still dominate public space and policy agendas [71]. Infrastructure quality plays a central role in enabling this shift: cyclists consistently prefer separation from motorized traffic, while poor pavement conditions discourage cycling [73]. Thus, the design, continuity, and inclusiveness of cycling infrastructure are key determinants of bikeability, requiring diverse and connected facilities that serve different users and trip purposes [72]. At the behavioral level, agent-based models simulate individual mobility decisions, showing that improvements in cycling safety or reductions in car comfort can drive meaningful shifts toward active transport [74]. Empirical studies indicate that many car trips could be replaced by low-carbon or active modes without significant increases in travel time [77, 78]. Recent computational frameworks use smartphone-based data to estimate modal-shift potential and emission reductions [75, 76]. However, most of these studies do not explicitly classify cycling infrastructure by safety level, nor do they analyze how upgrades between levels affect the modal shift from cars to bicycles.

Transport justice: Traditional transport models have mainly adopted a utilitarian perspective [79], aiming to maximize total welfare by improving aggregate outcomes, regardless of how benefits are distributed. In this view, inequalities are acceptable as long as they increase overall utility [80], which means that policies that disadvantage minorities or favor travelers with higher values of time can still be considered "optimal" [120]. As an alternative, Rawls' egalitarian theory [81, 121] and the capabilities approach have been proposed as more suitable foundations for transport justice, since they emphasize fair opportunities and attention to the least advantaged. Building on these ideas, Martens [7] introduces a sufficientarian perspective for transport, arguing that justice should focus on guaranteeing a minimum level of accessibility for everyone, while accessibility above this threshold may be left to market-based or self-financing mechanisms. This has led to the use of accessibility sufficiency indices as core indicators of justice in transport planning [122], and has been extended to multi-modal settings [9] and complemented by capabilities-based analyses across different contexts [82]. However, with the exception of [9, 10], most of this work remains largely evaluative: it identifies where and for whom accessibility is insufficient, but does not offer operational tools to design or optimize networks in line with these principles. To date, only [123] applies principles of justice to the optimization of a cycling network by targeting sufficient safety along paths, but it does not account for the impact on car travel times through potential reductions in road capacity. In conclusion, to the best of the authors' knowledge, there is still no prescriptive, operational framework that optimizes multi-modal (bike + car) networks in line with both utilitarian and sufficientarian principles, while simultaneously reducing cyclist safety risks and explicitly accounting for car travel-time accessibility impacts arising from road capacity reductions and induced mode shifts that modify origin–destination demand for both cyclists and drivers.

6.2 Model Design

6.2.1 Total system travel time

To assess accessibility and temporal efficiency in the cycling network, we convert the safety-adjusted cyclist distances into travel times using a representative average cycling speed $v_b > 0$. For an OD pair (u, v) and its (AoN) chosen bicycle path $p^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ under the upgraded configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, the corresponding bicycle path travel time is

$$T^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \sum_{a \in p^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}})} \frac{\bar{d}^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{v_b}. \tag{6.1}$$

Under AoN assignment, the bicycle total system travel time (TTT^{b}) is the demand-weighted sum of OD travel times,

$$\mathcal{TTT}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \sum_{(u,v) \in \mathcal{V}^{\mathrm{b}}_{\mathcal{O}} \times \mathcal{V}^{\mathrm{b}}_{\mathcal{D}}} \tilde{q}^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot T^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}). \tag{6.2}$$

Equivalently, by defining an arc travel time $\tau^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \bar{d}^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})/v_b$, the same quantity can be written in an arc-based form that mirrors the car definition:

$$\mathcal{TTT}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \tau^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \frac{\bar{d}^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{v_b}. \tag{6.3}$$

Finally, the total system travel time, combining car and bicycle components, is

$$\mathcal{TTT}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \mathcal{TTT}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + \mathcal{TTT}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}). \tag{6.4}$$

6.2.2 Interaction Between Cars and New Bike Lanes

As in Chapter 5, we model the creation (or upgrading) of cycling infrastructure as a potential reduction in car-lane capacity on shared road segments. Here, we adopt the same principle but refine the capacity-update rule to account explicitly for the number of car lanes and a road-width feasibility condition. Upgrading a bike lane may reduce the capacity of the adjacent car lane. However, as noted by [117], bike lanes that are not fully separated—specifically those with a DoS less than or equal to 4 do not significantly affect car travel times. In our framework, only upgrades to a DoS of at least 5 (i.e., fully separated infrastructure) are assumed to impact car traffic flow. The effective capacity of a car link, denoted $\tilde{\theta}^{\mathrm{c}}_a$, depends on its original capacity θ^{c}_a, the number of car lanes n^{c}_a, the road width l_a, and the upgraded bike lane DoS $\tilde{w}^{\mathrm{b}}_a$. To simplify notation, we introduce a binary indicator $z_a \in \{0, 1\}$ that equals 1 if the bike lane on link a is upgraded to a fully separated facility (i.e., $\tilde{w}^{\mathrm{b}}_a \geq 5$) on a narrow road ($l_a \leq 4.5$ m), and 0 otherwise. The effective capacity is then written compactly as

$$\tilde{\theta}^{\mathrm{c}}_a = \theta^{\mathrm{c}}_a \left(1 - \frac{z_a}{n^{\mathrm{c}}_a}\right). \tag{6.5}$$

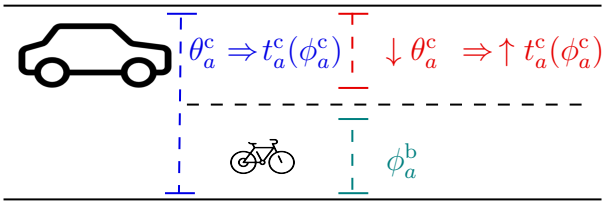

Figure 6.1: The creation/upgrading of a bike lane is modeled as reducing the capacity of the affected car link, resulting in an increased travel time $\uparrow t_a^c(\phi_a^c)$.

This formulation applies a capacity reduction only when a cycling upgrade is interpreted as a physical reallocation of roadway space to a protected facility, i.e., when the bike infrastructure reaches a fully separated standard (DoS ≥ 5) and the cross-section is narrow ($l_a \leq 4.5\,\text{m}$). The rationale is geometric: protected lanes or tracks typically require dedicated width and buffer/segregation elements, which in many street cross-sections can only be provided by removing (part of) a general-traffic lane or reallocating curb space. Design guidance for bicycle tracks and protected lanes reports typical widths that are comparable to a substantial fraction of a motor-vehicle lane [124, 125]; accordingly, we model this intervention as a reduction of effective car capacity. In contrast, lower DoS upgrades (e.g., DoS 2–3 paint-only or marking-based treatments) do not necessarily imply an immediate loss of car-lane capacity in a macroscopic model, since they can often be implemented within the existing carriageway (for instance through re-striping or advisory markings) and their geometric impact is more context-dependent. For this reason, we restrict the explicit capacity-loss mechanism to DoS ≥ 5, where a lane-reallocation interpretation is most defensible. Importantly, this does not mean that lower-DoS interventions have no impact on car traffic. Even painted lanes may influence behavior by improving cycling attractiveness and inducing mode shift or route changes. In our framework, such effects are represented through the demand side: cycling improvements modify generalized cycling cost and hence the bicycle mode share, which reduces car OD demand and can indirectly change congestion. Compared with the uniform reduction used in Chapter 5 (Section 5.2.3), the present formulation also captures that the marginal impact of removing one car lane decreases with the number of available lanes through the $\frac{1}{n_a^c}$ scaling factor. The threshold 4.5 m corresponds to standard lane-width guidance in France [126]. Fig. 6.1 illustrates the resulting interaction between capacity changes, congestion, and modal shift.

6.2.2.1 Mode Choice Model Based on Temporal Data

We develop a binary mode choice model linking changes in cycling network quality to variations in bicycle mode share over time, using historical data from the Grenoble metropolitan area. The annual cycling networks for 2015–2025 were extracted from OpenStreetMap using the `OSMnx` Python package [127]. For each year $k \in \{2015, 2016, \ldots, 2025\}$, we compute the

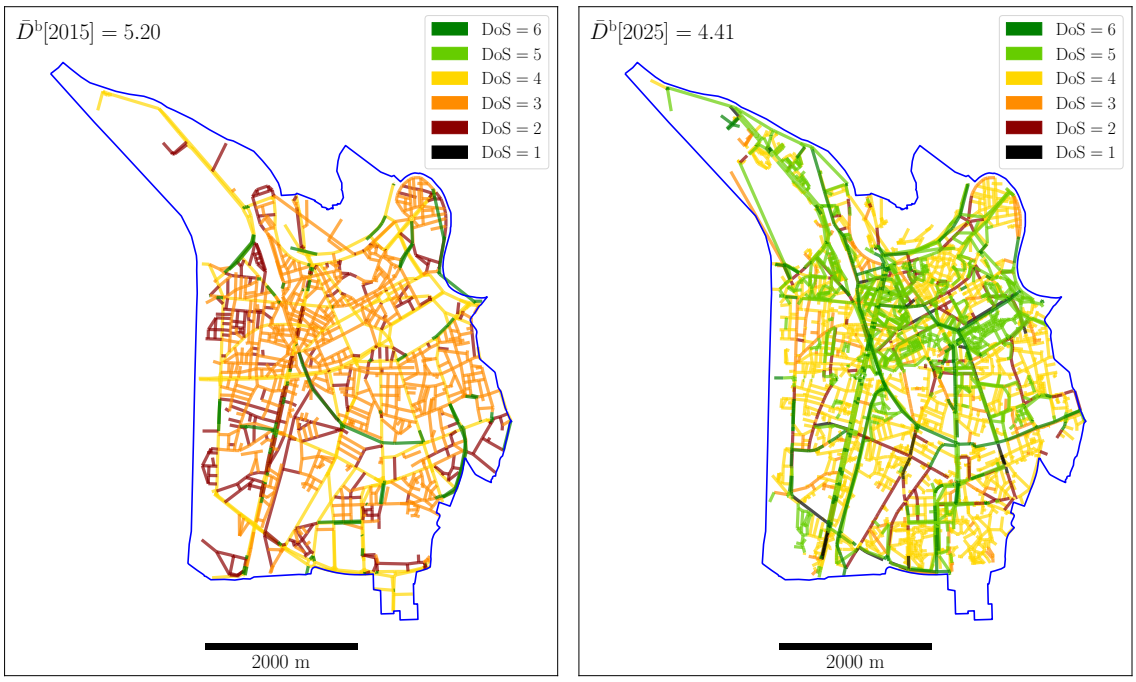

Figure 6.2: Historical cycling network of Grenoble, France. *Left:* year 2015. *Right:* year 2025. The average cycling distance $\overline{D}^{b}[k]$ decreases as the network expands and improves.

average (safety-adjusted) cycling distance as:

$$\overline{D}^{b}[k] = \frac{1}{|\mathcal{V}^{b}_{\mathcal{O}} \times \mathcal{V}^{b}_{\mathcal{D}}|} \sum_{(u,v) \in \mathcal{V}^{b}_{\mathcal{O}} \times \mathcal{V}^{b}_{\mathcal{D}}} \bar{d}^{b}_{(u,v)}[k], \tag{6.6}$$

where $\bar{d}^{b}_{(u,v)}[k]$ denotes the cyclist distance between origin u and destination v for the network configuration and traffic conditions of year k. Fig. 6.2 illustrates the Grenoble cycling network in 2015 and its projected configuration for 2025. Fig. 6.3 reports the average cycling distance and the bicycle mode share for the historical network between 2015 and 2025, based on observed values from [128–130], with intermediate years estimated through linear regression. As shown, there is an inverse relationship between the two variables: the average cycling distance decreases almost monotonically over time, while the bicycle mode share steadily increases, reflecting a progressive shift toward shorter but more frequent cycling trips as infrastructure expands and connectivity improves.

We acknowledge that cycling adoption is influenced by a wide range of factors, including social context, municipal initiatives, historical development, and gender dynamics [131, 132]. However, in this study we adopt a simplified formulation in which infrastructure is the sole explanatory variable of bicycle mode choice. This assumption is supported by evidence showing that the quality of cycling facilities strongly affects travel behavior, as cyclists consistently prefer routes separated from motorized traffic, while poor surface conditions tend to discourage cycling [73, 74]. We model the annual probability of choosing cycling, denoted as $P^{b}[k]$, based on the utility associated with the bicycle mode, $V^{b}[k]$. This utility is assumed

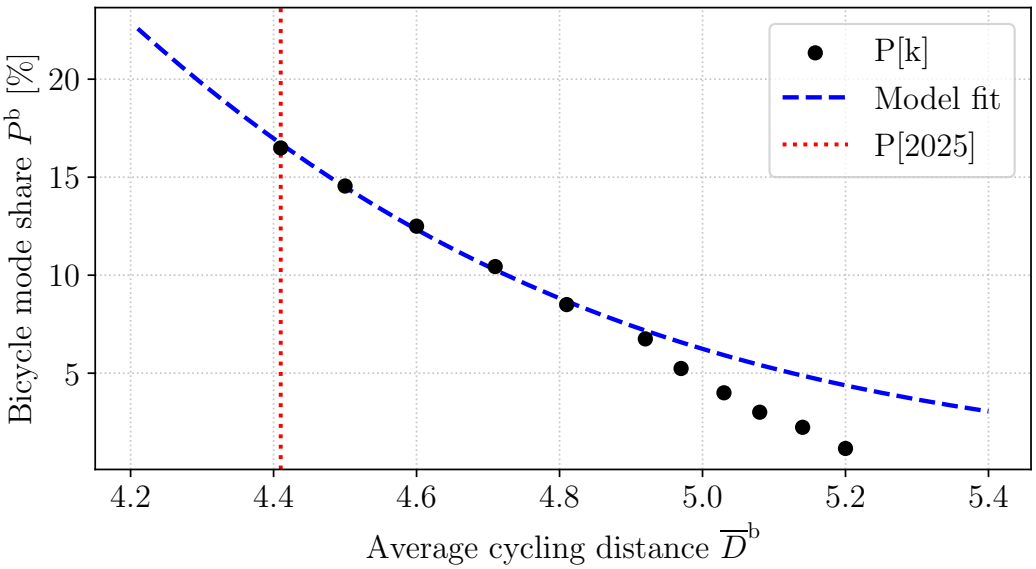

Figure 6.3: The logit model plot to predict $P^{b}[k]$ using $\bar{D}^{b}$.

to depend linearly on the average cycling distance $\overline{D}^{b}[k]$:

$$V^{b}[k] = \beta_0 + \beta_1 \overline{D}^{b}[k]. \tag{6.7}$$

Using a binary logit model, with the car mode as the reference alternative, the probability of choosing the bicycle mode in year k is expressed as:

$$P^{b}[k] = \frac{\exp(V^{b}[k])}{\exp(V^{b}[k]) + \exp(V^{c}[k])}. \tag{6.8}$$

Since only relative utilities influence the choice, we normalize the utility of the car mode to zero ($V^{c} \equiv 0$) without loss of generality. This yields the standard logit formulation:

$$P^{b}[k] = \frac{\exp(V^{b}[k])}{1 + \exp(V^{b}[k])} = \frac{1}{1 + \exp(-V^{b}[k])}, \tag{6.9}$$

where β_0 is an alternative-specific constant capturing unobserved factors favoring car use (e.g., travel time, fuel cost, or parking convenience), and $\beta_1 < 0$ represents the sensitivity of cycling adoption to the perceived distance or effort. Substituting Eq. (6.7) into the logit expression Eq. (6.9) gives the predicted bicycle mode share:

$$\widehat{P}^{b}[k] = \frac{1}{1 + \exp\big(-(\beta_0 + \beta_1 \overline{D}^{b}[k])\big)}. \tag{6.10}$$

The parameters (β_0, β_1) are estimated from the observed annual data pairs $\{(\overline{D}^{b}[k], P^{b}[k])\}$ by minimizing the squared difference between observed and predicted bicycle shares:

$$\min_{\beta_0,\beta_1} \sum_{k} \Big[P^{b}[k] - \widehat{P}^{b}[k]\Big]^2. \tag{6.11}$$

To ensure consistency with the current situation, β_0 is calibrated so that, for the existing 2025 network, the model reproduces the observed cycling share. Hence, if no network upgrades are applied (*i.e.*, $\tilde{\boldsymbol{w}}^{\mathrm{b}}$ identical to 2025), the predicted bicycle probability $\mathbb{P}^{\mathrm{b}}$ equals the observed value P^{b}_{2025}.

For any upgraded configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, the corresponding average cycling distance $\overline{D}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ determines the updated mode shares. Although $\overline{D}^{\mathrm{b}}[k]$ varies over time, it can also be viewed as a function of the network configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, since infrastructure upgrades that increase the DoS modify travel effort and route choices, thus affecting the resulting cycling share.

$$\mathbb{P}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \frac{1}{1+\exp\big(-(\beta_0+\beta_1\,\overline{D}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}))\big)}, \tag{6.12}$$

$$\mathbb{P}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = 1-\mathbb{P}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}), \tag{6.13}$$

and the total travel demand between each origin–destination pair (u,v) is reallocated as

$$\tilde{q}^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \mathbb{P}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\cdot Q_{(u,v)}, \tag{6.14}$$

$$\tilde{q}^{\mathrm{c}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \mathbb{P}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\cdot Q_{(u,v)}, \tag{6.15}$$

where $Q_{(u,v)}$ denotes the total travel demand of cars drivers and cyclists between origin u and destination v.

6.2.2.2 Surrogate-Based Flow Prediction with Feasibility Projection

Computing the UE traffic assignment as Problem 5.1 at each iteration is computationally expensive, particularly when travel demand and lane configurations change frequently due to network upgrades. To mitigate this cost, we employ a graph-based surrogate that predicts near-equilibrium car flows and travel times directly from network-level features, followed by a projection step to ensure feasibility.

Let $\boldsymbol{\phi}^{\mathrm{c}} = [\phi^{\mathrm{c}}_a]^\top_{a\in\mathcal{A}^{\mathrm{c}}} \in \mathbb{R}^{|\mathcal{A}^{\mathrm{c}}|}$ denote the vector of car flows on all arcs, and let $\boldsymbol{q}^{\mathrm{c}} = [q^{\mathrm{c}}_{(u,v)}]^\top_{(u,v)\in\mathcal{V}^{\mathrm{b}}_{\mathcal{O}}\times\mathcal{V}^{\mathrm{b}}_{\mathcal{D}}} \in \mathbb{R}^{|\mathcal{V}^{\mathrm{c}}_{\mathcal{O}}\times\mathcal{V}^{\mathrm{c}}_{\mathcal{D}}|}$ represent the OD demand vector. The UE flow $\boldsymbol{\phi}^{\mathrm{c}^\star}$ satisfies the flow-conservation constraint

$$B\boldsymbol{\phi}^{\mathrm{c}^*} = \boldsymbol{d}^{\mathrm{c}}, \tag{6.16}$$

where $B \in \{-1,0,1\}^{|\mathcal{V}^{\mathrm{c}}|\times|\mathcal{A}^{\mathrm{c}}|}$ is the node–arc incidence matrix and $\boldsymbol{d}^{\mathrm{c}} \in \mathbb{R}^{|\mathcal{V}^{\mathrm{c}}|}$ is the net demand vector (positive at origins, negative at destinations, and zero at intersections and junctions).

We train a Graph Convolutional Network (GCN) surrogate $\mathcal{M}$ on UE traffic assignments obtained under random bike-lane upgrade configurations $\tilde{\boldsymbol{w}}^{\mathrm{b}}$. Each configuration alters the car network capacities and OD demands, generating distinct equilibrium flows and travel times used as ground truth. The GCN takes as input the car network topology $\mathcal{A}^{\mathrm{c}}$, the effective arc capacities $\tilde{\boldsymbol{\theta}}^{\mathrm{c}} = [\tilde{\theta}^{\mathrm{c}}_a]_{a\in\mathcal{A}^{\mathrm{c}}}$ from Eq. (6.5), and the post-shift OD demand vector

$\tilde{\boldsymbol{q}}^{\mathrm{c}} = [\tilde{q}^{\mathrm{c}}_{(u,v)}]^{\top}_{(u,v)\in\mathcal{V}^{\mathrm{c}}_{\mathcal{O}}\times\mathcal{V}^{\mathrm{c}}_{\mathcal{D}}}$ from Eq. (6.14), and learns the mapping from network configuration to predict UE traffic flow distribution $\tilde{\boldsymbol{\phi}}^{\mathrm{c}}$:

$$\mathcal{M}:\ (\mathcal{A}^{\mathrm{c}},\ \tilde{\boldsymbol{\theta}}^{\mathrm{c}},\ \tilde{\boldsymbol{q}}^{\mathrm{c}})\ \longmapsto\ \tilde{\boldsymbol{\phi}}^{\mathrm{c}}. \tag{6.17}$$

Since these inputs, in particular $\tilde{\boldsymbol{\theta}}^{\mathrm{c}}$ and potentially $\tilde{\boldsymbol{q}}^{\mathrm{c}}$ through mode choice are functions of the upgraded cycling configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, we write the surrogate compactly as

$$\mathcal{M}:\ \tilde{\boldsymbol{w}}^{\mathrm{b}}\ \longmapsto\ \tilde{\boldsymbol{\phi}}^{\mathrm{c}}. \tag{6.18}$$

However, since the GCN does not explicitly enforce flow conservation, $\tilde{\boldsymbol{\phi}}^{\mathrm{c}}$ may violate $B\tilde{\boldsymbol{\phi}}^{\mathrm{c}} = \boldsymbol{d}^{\mathrm{c}}$. We restore feasibility by projecting the surrogate prediction onto the flow-conserving subspace via the quadratic program

$$\tilde{\boldsymbol{\phi}}^{c^\star}\ =\ \arg\min_{\boldsymbol{\phi}^{\mathrm{c}}\geq 0}\ \left\|\boldsymbol{\phi}^{\mathrm{c}} - \tilde{\boldsymbol{\phi}}^{\mathrm{c}}\right\|_2^2\quad \text{s.t.}\quad B\boldsymbol{\phi}^{\mathrm{c}} = \boldsymbol{d}^{\mathrm{c}}. \tag{6.19}$$

The solution $\tilde{\boldsymbol{\phi}}^{c^\star}$ represents the Euclidean projection of the GCN-predicted flows onto the feasible subspace, ensuring node-level flow conservation and nonnegativity. The corresponding car total system travel time associated with $\tilde{\boldsymbol{\phi}}^{c^\star}$ is also denoted as $\mathcal{TT}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$. Finally, the surrogate-based total system travel time combines the computed bicycle component with the GCN-predicted car component:

$$\mathcal{TT}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\ =\ \mathcal{TT}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\ +\ \mathcal{TT}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}). \tag{6.20}$$

This GCN–projection framework provides fast and consistent approximations of network performance, enabling efficient optimization of multi-modal configurations while substantially reducing computational cost. A detailed assessment of the model's predictive accuracy and reliability is presented in the Results section.

6.2.3 Optimization Problem (Utilitarian Perspective)

We now introduce the utilitarian formulation of the trade-off problem for upgrading cycling infrastructure. Given a fixed budget B, the planner selects a subset of road segments to upgrade in order to enhance their degree of separation. The objective balances two competing goals: (i) minimizing cyclists' average exposure to motor traffic, and (ii) limiting the aggregate increase in total travel times due to car capacity reallocation. This formulation embodies a utilitarian principle, wherein the planner aims to minimize the overall system cost by jointly reducing traffic exposure and travel-time inefficiency across all network users. Let the decision variables

$$\tilde{\boldsymbol{w}}^{\mathrm{b}} = \left[\tilde{w}^{\mathrm{b}}_a\right]^{\top}_{a\in\mathcal{A}^{\mathrm{b}}} \tag{6.21}$$

denote the vector of upgraded DoS values for all bicycle arcs, with feasible set

$$\Omega = \left\{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \{1,\dots,6\}^{|\mathcal{A}^{\mathrm{b}}|} : w^{\mathrm{b}}_a \leq \tilde{w}^{\mathrm{b}}_a \leq 6,\ \forall a \in \mathcal{A}^{\mathrm{b}}\right\}. \tag{6.22}$$

Table 6.1 summarizes the main quantities used in the optimization model. We formulate the

Symbol	Description	Domain	Unit
$\phi_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Bicycle flow on arc a	$\mathbb{R}_{\geq 0}$	bike/hr
d_a^{b}	Length of bike arc a	$\mathbb{R}_{>0}$	m
$\bar{d}_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Cyclist distance on arc a	$\mathbb{R}_{>0}$	m
$\mathcal{E}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Cyclist exposure on arc a	$\mathbb{R}_{\geq 0}$	–
$\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$	Upgrade cost of arc a	$\mathbb{R}_{\geq 0}$	€
$\tilde{\theta}_a^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Updated car lane capacity on arc a	$\mathbb{R}_{\geq 0}$	veh/hr
$\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Predicted total travel time	$\mathbb{R}_{\geq 0}$	hrs

Table 6.1: Notation used in the utilitarian optimization formulation.

problem as a bi-level optimization, where the upper level determines infrastructure upgrades, and the lower level models the resulting cyclist route choices.

At the *upper level* ($\mathbb{U}_{\mathrm{pper}}$), the planner selects upgrades $\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega$ subject to the budget B, minimizing total cyclist exposure and total time while controlling the impact on the average car total travel times. The formulation naturally prioritizes the improvement of longer links with both high cyclist and vehicle volumes, since these contribute most to total exposure and travel cost.

At the *lower level* ($\mathbb{L}_{\mathrm{ower}}$), given a specific upgraded configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$, cyclist flows are assigned to origin–destination pairs using an all-or-nothing (AoN) shortest-path assignment based on the effective cyclist distance defined in Section 4.2.1. This ensures that the induced flow distribution reflects perceived safety across the network.

Problem 6.1 (Utilitarian Trade-Off Network Upgrade).

$\mathbb{U}_{\mathrm{pper}}$:

$$\boldsymbol{w}^{\mathrm{b}\star} = \arg\min_{\tilde{\boldsymbol{w}}^{\mathrm{b}} \in \Omega} \quad \gamma \cdot \underbrace{\sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot f_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot d_a^{\mathrm{b}}}_{\text{Cyclist exposure } \mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}})} + (1-\gamma) \cdot \underbrace{\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}_{\text{Predicted total travel time}} \tag{6.23}$$

$$\text{s.t.} \quad \sum_{a \in \mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \leq B, \tag{6.24}$$

$$\tilde{\theta}_a^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \theta_a^{\mathrm{c}} \cdot \left(1 - \frac{z_a}{n_a^{\mathrm{c}}}\right), \tag{6.25}$$

$\mathbb{L}_{\mathrm{ower}}$:

$$\boldsymbol{\phi}^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \arg\min_{\boldsymbol{\phi}^{\mathrm{b}} \in \mathbb{R}_{\geq 0}^{|\mathcal{A}^{\mathrm{b}}|}} \sum_{a \in \mathcal{A}^{\mathrm{b}}} \bar{d}_a^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \phi_a^{\mathrm{b}} \tag{6.26}$$

$$\text{s.t.} \quad \phi_a^{\mathrm{b}} = \sum_{(u,v) \in \mathcal{V}_{\mathcal{O}}^{\mathrm{b}} \times \mathcal{V}_{\mathcal{D}}^{\mathrm{b}}} \tilde{q}_{(u,v)}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \cdot \mathbb{I}_{a \in p_{(u,v)}^{\mathrm{b}}}, \quad \forall a \in \mathcal{A}^{\mathrm{b}}.$$

Here, $\gamma \in [0,1]$ balances the trade-off between reducing cyclist exposure and minimizing

system total travel time. The budget constraint Eq. (6.24) ensures feasibility of investments, while capacity reductions Eq. (6.25) reflect lane reallocation when high-quality (DoS ≥ 5) bike lanes are introduced. The lower-level problem Eq. (6.26) assigns bicycle flows via an all-or-nothing shortest path model under updated cyclist distances and the new bike demands following mode shifts. These optimal flows $\boldsymbol{\phi}^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ feed back into the upper-level objective, ensuring consistency between infrastructure upgrades, flows, and exposures. Problem 6.1 is NP-hard. This can be deduced by reduction from the 0–1 knapsack problem via a restricted instance of our formulation. Consider (i) setting $\gamma = 1$ so that the travel-time term is removed (or, equivalently, assuming $\tilde{\theta}_a^{\mathrm{c}} = \theta_a^{\mathrm{c}}$ for all a), and (ii) fixing the lower-level bicycle assignment so that ϕ_a^{b} is constant (e.g., equal to its baseline value). Then the upper-level decision reduces to selecting a subset of arc upgrades under the budget constraint $\sum_{a\in\mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \leq B$ to maximize a fixed exposure reduction, which is equivalent to a 0–1 knapsack problem (or a multi-choice knapsack if multiple upgrade levels per arc are allowed). Since this restricted problem is NP-hard, the full bi-level problem—which additionally couples upgrades with route choice and includes nonlinear effects through $\tilde{\theta}_a^{\mathrm{c}}$ is also NP-hard. See [53] for a formal statement in a closely related setting.

To clarify, we emphasize that car traffic is not held constant in Problem 6.1. The model includes an induced-demand (mode-choice) mechanism (Section 6.2.2.1) that reallocates the total OD demand $Q_{(u,v)}$ between cars and bicycles as a function of the upgraded cycling configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$. Hence both $\tilde{\boldsymbol{q}}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ and $\tilde{\boldsymbol{q}}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ vary with the upgrades. Given the updated car demand and the capacity changes induced by separated cycling facilities Eq. (6.25), car flows are then updated through the surrogate-based traffic assignment described in Section 6.2.2.2: the GCN predicts a near-UE flow pattern $\tilde{\boldsymbol{\phi}}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$, which is subsequently projected onto the flow-conservation constraints to obtain a feasible approximation $\tilde{\boldsymbol{\phi}}^{\mathrm{c}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$. The resulting predicted car travel-time term $\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{c}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ is therefore upgrade-dependent. Finally, the total system travel time in the objective, $\mathcal{T}\tilde{\mathcal{T}}\mathcal{T}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$, combines the bicycle component (which changes through $\vec{d}_a^{\mathrm{b}}$ and rerouted bicycle flows) with this surrogate-based car component (Eq. (6.20)). In other words, the framework captures the primary objective of cycling upgrades inducing demand shifts from cars to bikes while replacing repeated exact UE solves with a surrogate model.

6.2.4 Optimization Problem (Sufficientarian Perspective)

Now, we incorporate a sufficientarian perspective into Problem 6.1 by ensuring that as many individuals as possible achieve a minimum level of safety on paths throughout the network. Using the a similar definition of Eq. (4.5), for a path $p_{(u,v)}^{\mathrm{b}}$ between OD pair (u, v), we define the safety score as

$$S_{p_{(u,v)}^{\mathrm{b}}} = \frac{1}{\sum_{a\in p_{(u,v)}^{\mathrm{b}}} \phi_a^{\mathrm{b}} \cdot f_a^{\mathrm{b}} \cdot d_a^{\mathrm{b}} + \epsilon}, \tag{6.27}$$

where ϕ_a^{b} is the cyclist flow on arc a, f_a^{b} is the traffic-related risk factor, d_a^{b} is its length, and $\epsilon > 0$ avoids division by zero and limits the maximum safety score. This score reflects the inverse relationship between cyclist exposure and perceived safety, consistent with established

findings in route choice modeling [133], where cyclists prefer routes with lower traffic volumes. Although this formulation is a simplification, it serves as a continuous proxy for perceived safety when only exposure data is available. To determine what constitutes a sufficient level of safety, we adopt a relative threshold approach. Let the global average safety across all regions be

$$S_{\mathcal{G}}^{\mathrm{b}} = \frac{1}{|\mathcal{V}^{\mathrm{b}}|^2} \sum_{(u,v)\in\mathcal{V}^{\mathrm{b}}\times\mathcal{V}^{\mathrm{b}}} S_{p_{(u,v)}^{\mathrm{b}}}. \tag{6.28}$$

We then define a relative path-safety threshold s_{thr} as a proportion of the average safety level $s_{\mathcal{G}}^{\mathrm{b}}$, using a parameter λ that takes values from the set $\{10\%, 20\%, 30\%, 40\%, 50\%\}$.

$$S_{\mathrm{thr}} = \lambda S_{\mathcal{G}}^{\mathrm{b}}. \tag{6.29}$$

This method follows the standard practice proposed by Martens in the context of transport justice [7, 122], where sufficiency thresholds are defined as a proportion of an average (or median) accessibility level. Cycle paths whose safety scores fall below this threshold ($S_{p_{(u,v)}^{\mathrm{b}}} \leq S_{\mathrm{thr}}$) are considered insufficiently protected. To quantify this shortfall, and following a similar approach to [9], we adopt a linear form of the Foster–Greer–Thorbecke (FGT) index, originally developed for poverty measurement [134]. In the FGT framework, deprivation is measured as the normalized gap between an individual's income and a predefined poverty line; analogously, here we measure the safety insufficiency of each path as the normalized deficit of its safety score relative to the threshold

$$\mathcal{I}_{p_{(u,v)}^{\mathrm{b}}} = \max\left\{0,\ \frac{S_{\mathrm{thr}} - S_{p_{(u,v)}^{\mathrm{b}}}}{S_{\mathrm{thr}}}\right\}. \tag{6.30}$$

This formulation provides a continuous and interpretable measure of how far below the sufficiency threshold a given cycling path falls, enabling aggregate indicators of safety equity across the network. Dividing the network into regions $\mathcal{R} = \{1, \ldots, R\}$, where $\mathcal{V}_r^{\mathrm{b}} \subseteq \mathcal{V}^{\mathrm{b}}$ represents nodes in region r, allows for regional analysis, then, the regional safety insufficiency is computed as the demand-weighted average safety insufficiency for each OD pair within the region

$$\mathcal{I}_r(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \frac{\sum_{u\in\mathcal{V}_r^{\mathrm{b}}} \mathcal{I}_{p_{(u,v)}^{\mathrm{b}}} \cdot \tilde{q}_{(u,v)}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{\sum_{u\in\mathcal{V}_r^{\mathrm{b}}} \tilde{q}_{(u,v)}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}, \quad \forall v \in \mathcal{V}^{\mathrm{b}}, \tag{6.31}$$

where $q_{(u,v)}^{\mathrm{b}}$ is the demand for OD pair (u, v). The total population-weighted safety insufficiency across the network is given by

$$J_{\mathcal{SI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \frac{\sum_{r\in\mathcal{R}} n_r^{\mathrm{b}}\, \mathcal{I}_r(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{\sum_{r\in\mathcal{R}} n_r^{\mathrm{b}}}. \tag{6.32}$$

where n_r represents the number of cyclists in region r.

To formulate the sufficientarian perspective in terms of temporal accessibility for both cars and bicycles, we first introduce a path travel time threshold denoted by T_{thr} (e.g., $T_{\mathrm{thr}} = 20\ \mathrm{min} = 1200\,\mathrm{s}$, as used in [9]). For bicycles, given any trip (u, v) and its corresponding

chosen path $p_{(u,v)}$, and using the previously defined path travel time $T^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}}$ in Eq. (6.1), the unaccessibility is defined as

$$\mathcal{I}^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}} \;=\; \max\left\{0,\; \frac{T^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}} - T_{\mathrm{thr}}}{T_{\mathrm{thr}}}\right\}. \tag{6.33}$$

This quantity measures how much longer a bicycle trip takes compared to the acceptable travel-time threshold T_{thr}. A trip between u and v is considered accessible if its travel time does not exceed the threshold ($T^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}} \leq T_{\mathrm{thr}}$), and unaccessible otherwise. Hence, $\mathcal{I}^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}}$ represents the relative excess of travel time beyond this limit. To capture cyclists accessibility unfairness at a regional level, we consider the demand-weighted average of trip-level unaccessibilities for all trips originating in region r:

$$\mathcal{I}^{\mathrm{b}}_{r}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \frac{\sum_{(u,v):\, u\in\mathcal{V}^{\mathrm{b}}_{r}} \mathcal{I}^{\mathrm{b}}_{p^{\mathrm{b}}_{(u,v)}} \cdot \tilde{q}^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{\sum_{(u,v):\, u\in\mathcal{V}^{\mathrm{b}}_{r}} \tilde{q}^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}. \tag{6.34}$$

This formulation ensures that regions with higher cycling demand exert a proportionally greater influence on the overall accessibility measure. To evaluate cyclist accessibility unfairness at the network level, we define the cyclist-weighted network time accessibility unfairness, which aggregates regional accessibilities unfairness according to the number of cyclists n^{b}_{r} in each region:

$$J^{\mathrm{b}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \;=\; \frac{\sum_{r\in\mathcal{R}} n^{\mathrm{b}}_{r} \cdot \mathcal{I}^{\mathrm{b}}_{r}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}{\sum_{r\in\mathcal{R}} n^{\mathrm{b}}_{r}}. \tag{6.35}$$

This indicator provides a system-level measure of time accessibility unfairness, reflecting the share of the cycling population whose trips exceed acceptable travel-time limits. Analogously, we define drivers-weighted network time accessibility unfairness $J^{\mathrm{c}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$ for cars.

Finally, we define the sufficientarian bi-level optimization problem, where $\mathbb{U}_{\mathrm{pper}}$ ensures that as many cyclists as possible reach a path-safety threshold and as many cyclist and drivers as possible meet the minimum path travel time threshold. As in the utilitarian formulation, the lower level $\mathbb{L}_{\mathrm{ower}}$ determines the resulting cyclist flow distribution under the upgraded network.

Problem 6.2 (Sufficientarian Trade-Off Network-Upgrade).

$$\mathbb{U}\text{pper}:$$

$$\tilde{\boldsymbol{w}}^{\mathrm{b}\star} = \arg\min_{\tilde{\boldsymbol{w}}^{\mathrm{b}}\in\Omega} \quad \gamma \cdot \underbrace{J_{\mathcal{SI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})}_{\text{Safety insufficiency}} + (1-\gamma)\cdot \underbrace{\left[J^{\mathrm{b}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) + J^{\mathrm{c}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\right]}_{\text{Accessibility unfairness}} \tag{6.36}$$

$$\text{s.t.} \sum_{a\in\mathcal{A}^{\mathrm{b}}} d^{\mathrm{b}}_a \cdot \pi(w^{\mathrm{b}}_a, \tilde{w}^{\mathrm{b}}_a) \;\leq\; B, \tag{6.37}$$

$$\tilde{\theta}^{\mathrm{c}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \theta^{\mathrm{c}}_a \cdot \left(1 - \frac{z_a}{n^{\mathrm{c}}_a}\right), \tag{6.38}$$

$$\mathbb{L}\text{ower}:$$

$$\boldsymbol{\phi}^{\mathrm{b}\star}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) = \arg\min_{\boldsymbol{\phi}^{\mathrm{b}}\in\mathbb{R}^{|\mathcal{A}^{\mathrm{b}}|}_{\geq 0}} \sum_{a\in\mathcal{A}^{\mathrm{b}}} \bar{d}^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})\cdot\phi^{\mathrm{b}}_a \tag{6.39}$$

$$\text{s.t.} \quad \phi^{\mathrm{b}}_a = \sum_{(u,v)\in\mathcal{V}^{\mathrm{b}}_{\mathcal{O}}\times\mathcal{V}^{\mathrm{b}}_{\mathcal{D}}} \tilde{q}^{\mathrm{b}}_{(u,v)}(\tilde{\boldsymbol{w}}^{\mathrm{b}})\cdot\mathbb{I}_{a\in p^{\mathrm{b}}_{(u,v)}}, \quad \forall a\in\mathcal{A}^{\mathrm{b}}.$$

A summary of the model notation is given in Table 6.2.

Symbol	Description	Domain	Unit
$q^{b}_{(u,v)}$	Bike demand from u to v	$\mathbb{R}_{\geq 0}$	bike/hr
$q^{c}_{(u,v)}$	Car demand from u to v	$\mathbb{R}_{\geq 0}$	veh/hr
$\bar{d}^{\mathrm{b}}_{(u,v)}$	Cyclist distance u-v	$\mathbb{R}_{\geq 0}$	m
AADT(a)	AADT adjacent to $a\in\mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	veh/hr
S_{thr}	Safety path threshold	$\mathbb{R}_{\geq 0}$	-
T_{thr}	Time path threshold	$\mathbb{R}_{\geq 0}$	-
$\phi^{\mathrm{b}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Bicycle flow on arc $a\in\mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	bike/hr
$\phi^{\mathrm{c}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Car flow on arc $a\in\mathcal{A}^{\mathrm{c}}$	$\mathbb{R}_{\geq 0}$	car/hr
$\bar{D}^{\mathrm{b}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Average cyclist distance	$\mathbb{R}_{\geq 0}$	m
$\mathcal{E}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Cyclist exposure on $a\in\mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	-
$\tilde{\theta}^{\mathrm{c}}_a(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Updated car capacity on $a\in\mathcal{A}^{\mathrm{c}}$	$\mathbb{R}_{\geq 0}$	veh/hr
$\pi(w^{\mathrm{b}}_a, \tilde{w}^{\mathrm{b}}_a)$	Upgrade cost on $a\in\mathcal{A}^{\mathrm{b}}$	$\mathbb{R}_{\geq 0}$	€
$\mathcal{TTT}^{\mathrm{tot}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Predicted total system travel time	$\mathbb{R}_{\geq 0}$	-
$\mathcal{E}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Total cyclist exposure	$\mathbb{R}_{\geq 0}$	-
$J_{\mathcal{SI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Total safety insufficiency	$\mathbb{R}_{\geq 0}$	-
$J^{\mathrm{b}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Total bike-time insufficiency	$\mathbb{R}_{\geq 0}$	-
$J^{\mathrm{c}}_{\mathcal{TI}}(\tilde{\boldsymbol{w}}^{\mathrm{b}})$	Total car-time insufficiency	$\mathbb{R}_{\geq 0}$	-

Table 6.2: Abbreviated notation for variables in trade-off network model. Notation for optimization variables depending on the upgrade configuration $\tilde{\boldsymbol{w}}^{\mathrm{b}}$

Now we will introduce the case study of this work.

6.3 Case Study

This section introduces a case study of Grenoble, France. It also examines the results obtained using the MLCC, GA-U and GA-S algorithms in comparison to the original scenario.

6.3.1 Algorithm Comparisons

Using the historical data from Fig 6.3, the estimated parameters of the bicycle mode-choice model are $\beta_0 = 0.128$ and $\beta_1 = -0.886$, as illustrated in Fig. 6.3, in this logit formulation, β_0 is an alternative-specific constant that captures unobserved factors favoring cycling over car use—such as cultural preferences, infrastructure quality, or policy incentives—when the perceived cycling distance is zero. The parameter $\beta_1 < 0$ quantifies the sensitivity of cycling adoption to the average cycling distance: as distances increase, the perceived disutility of cycling rises, leading to a lower probability of choosing the bicycle mode. The estimated values indicate a moderate baseline preference for cycling and a strong negative influence of distance on mode choice, consistent with observed behavioral patterns in the Grenoble metropolitan area. The generation of 912 samples for the UE required 5,928 CPU hours. Fig. 6.4(a)–(b) illustrate the evolution of the training loss, testing loss, and their generalization gap over 800 epochs, obtained using PyTorch [135] to train the GCN. Here, the loss is measured in mean squared error (MSE), and the gap corresponds to $\mathrm{MSE}_{\mathrm{test}} - \mathrm{MSE}_{\mathrm{train}}$, which provides a simple indicator of overfitting. Fig. 6.4(b) further compares the surrogate predictions $\tilde{\phi}_a^{\mathrm{c}}$ against the UE reference flows $\phi_a^{\mathrm{c}^\star}$ on each car arc a. Fig. 6.4(c) reports the residuals $\tilde{\phi}_a^{\mathrm{c}} - \phi_a^{\mathrm{c}^\star}$ as a function of the UE reference flow $\phi_a^{\mathrm{c}^\star}$. The residuals are centered around zero across the full flow range, indicating no strong systematic over or underestimation. As expected, the spread of residuals increases with flow magnitude (heteroscedasticity), meaning that absolute errors are larger on high-flow links even though relative errors remain moderate.

To substantiate the approximation quality, we report standard regression error metrics between predicted and UE flows. The mean absolute error (MAE) is 68.9 veh/h and the root mean squared error (RMSE) is 105.9 veh/h. The coefficient of determination is $R^2 = 0.9696$, indicating that the surrogate explains most of the variance in UE flows. In addition, 51.5% of link flows are predicted within ±5% of the UE value, and 82.1% within ±10%. Finally, the 95th and 99th percentiles of the absolute error are 227 and 369 veh/h, respectively, showing that the largest discrepancies occur on a small fraction of high-flow links. Overall, these results indicate that the GCN provides a reasonably accurate and computationally efficient approximation of UE flows, suitable for use as a surrogate within the optimization loop, while still motivating the feasibility-projection step introduced in Section 6.2.2.2.

As in the previous chapthers, let $B_{\max}$ represent the budget required to upgrade the entire network to DoS = 6. The cost per meter $\pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$ of upgrading arc a from w_a^{b} to $\tilde{w}_a^{\mathrm{b}}$ is given by the matrix below. The total cost of upgrading an arc is computed as $d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}})$, where d_a^{b} is the arc length, making the cost distance-dependent. The cost matrix Π is adapted from Cerema's reported unit construction costs for cycling infrastructure upgrades in France [110] in Eq. 3.43.

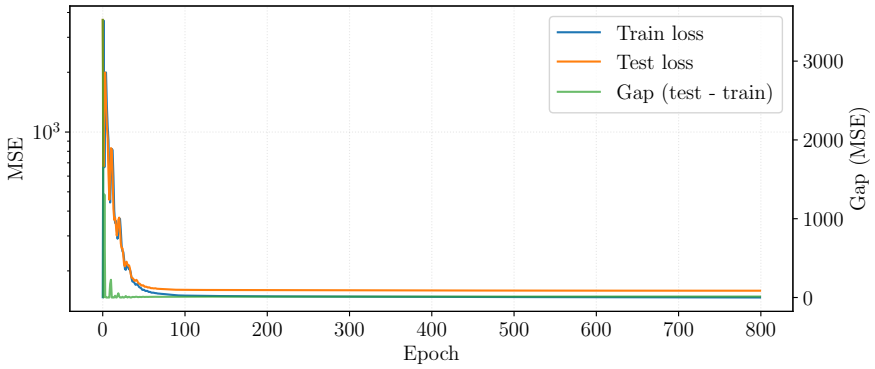

(a) Training, test and gap progress.

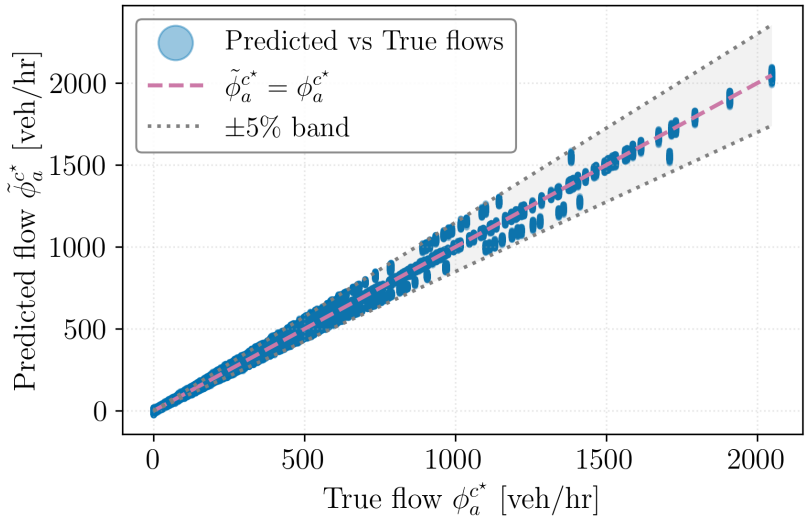

(b) Predicted vs true flows using the surrogate model

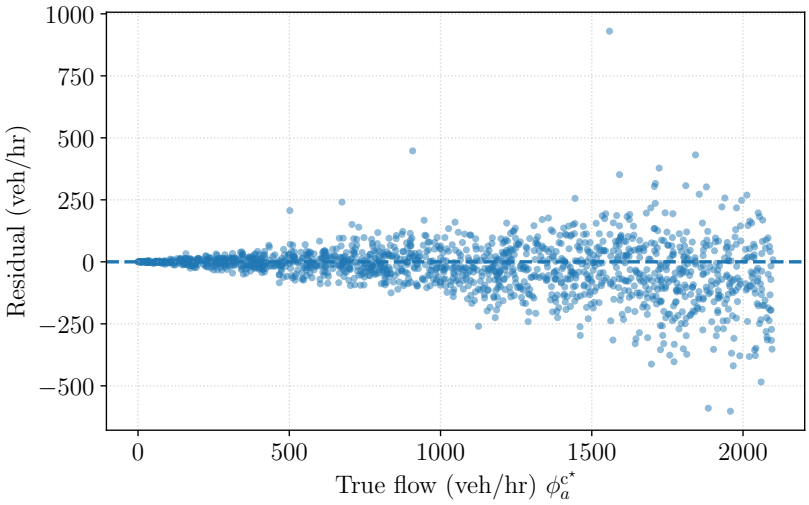

(c) Residuals vs True flow

Figure 6.4: On the top we present the epochs evolution with the 912 samples created for solving Problem 5.1. On the middle how the predicted flows approaches the true flow found by the UE, and on the. And on the bottom we present the residual analysis.

According to the Grenoble cycling investment plan (March 2025) [136], the planned budget is about 24.5M €, which corresponds to $B \approx 0.15\, B_{\max}$. We therefore solve Problem 6.1 (utilitarian) and Problem 6.2 (sufficientarian) under this budget constraint. For each formulation, we explore the trade-off between the competing objectives by solving the corresponding bi-level problem for a discrete set of weights

$$\Lambda = \left\{0, \tfrac{1}{10}, \tfrac{1}{5}, \tfrac{3}{10}, \tfrac{2}{5}, \tfrac{1}{2}, \tfrac{3}{5}, \tfrac{7}{10}, \tfrac{4}{5}, \tfrac{9}{10}, 1\right\},$$

where $\gamma \in \Lambda$ controls the balance between the safety-related term and the travel-time/accessibility-related term. This yields, for each γ, one candidate solution and one point on the corresponding Pareto front (reported in Fig. 6.5). In the sufficientarian formulation, we additionally test several safety-threshold levels through

$$\lambda \in \left\{\tfrac{1}{10}, \tfrac{1}{5}, \tfrac{3}{10}, \tfrac{2}{5}, \tfrac{1}{2}\right\},$$

which define increasingly demanding path-safety requirements. For each λ, the algorithm is run for all $\gamma \in \Lambda$, producing a family of trade-off curves.

Because the two objectives are expressed in different units, we do not define an "optimal" compromise point (e.g., via Euclidean distance to the origin). Instead, selecting γ is ultimately a policy decision that reflects how decision makers value safety improvements relative to travel-time impacts. To support interpretation and visualization of the results, we highlight one representative weight $\gamma^\star_\lambda$ for each threshold level λ, chosen solely for illustration:

$$\gamma^\star_{\frac{1}{10}} = \tfrac{1}{5},\ \gamma^\star_{\frac{1}{5}} = \tfrac{3}{10},\ \gamma^\star_{\frac{3}{10}} = \tfrac{4}{10},\ \gamma^\star_{\frac{2}{5}} = \tfrac{5}{10},\ \gamma^\star_{\frac{1}{2}} = \tfrac{6}{10}.$$

These $\gamma^\star_\lambda$ values are used only to exemplify the behavior of the utilitarian and sufficientarian solutions in subsequent figures and discussion; they should not be interpreted as mathematically optimal trade-offs.

Figure 6.5 compares the sets of solutions produced by GA-U (Problem 6.1) and GA-S (Problem 6.1). As expected, small γ values emphasize minimizing travel-time/accessibility impacts, while large γ values emphasize cyclist safety objectives, typically at the cost of higher travel-time impacts. Lower safety-threshold levels (smaller λ) generally allow larger safety gains with smaller impacts on car travel times, whereas higher λ values impose stricter requirements and may lead to larger disruptions in the motorized network. The figure also reveals a difference in convergence behavior between the two formulations. In this case, the sufficientarian approach appears to converge faster than the utilitarian one. A possible explanation is that the sufficientarian objective concentrates the search on the subset of regions and cyclists that remain below the prescribed safety threshold. Once a substantial proportion of the network already satisfies this requirement, the optimization pressure becomes focused on a smaller set of critical deficits, which can make the search direction more explicit and reduce the number of relevant competing alternatives. By contrast, the utilitarian formulation rewards aggregate improvement over the entire network, including areas that are already relatively well performing. As a result, many distinct intervention patterns may yield similar gains in the objective, which broadens the set of attractive solutions and can delay stabilization. In this sense, the faster convergence of the sufficientarian formulation may stem from the more selective and targeted nature of its objective function.

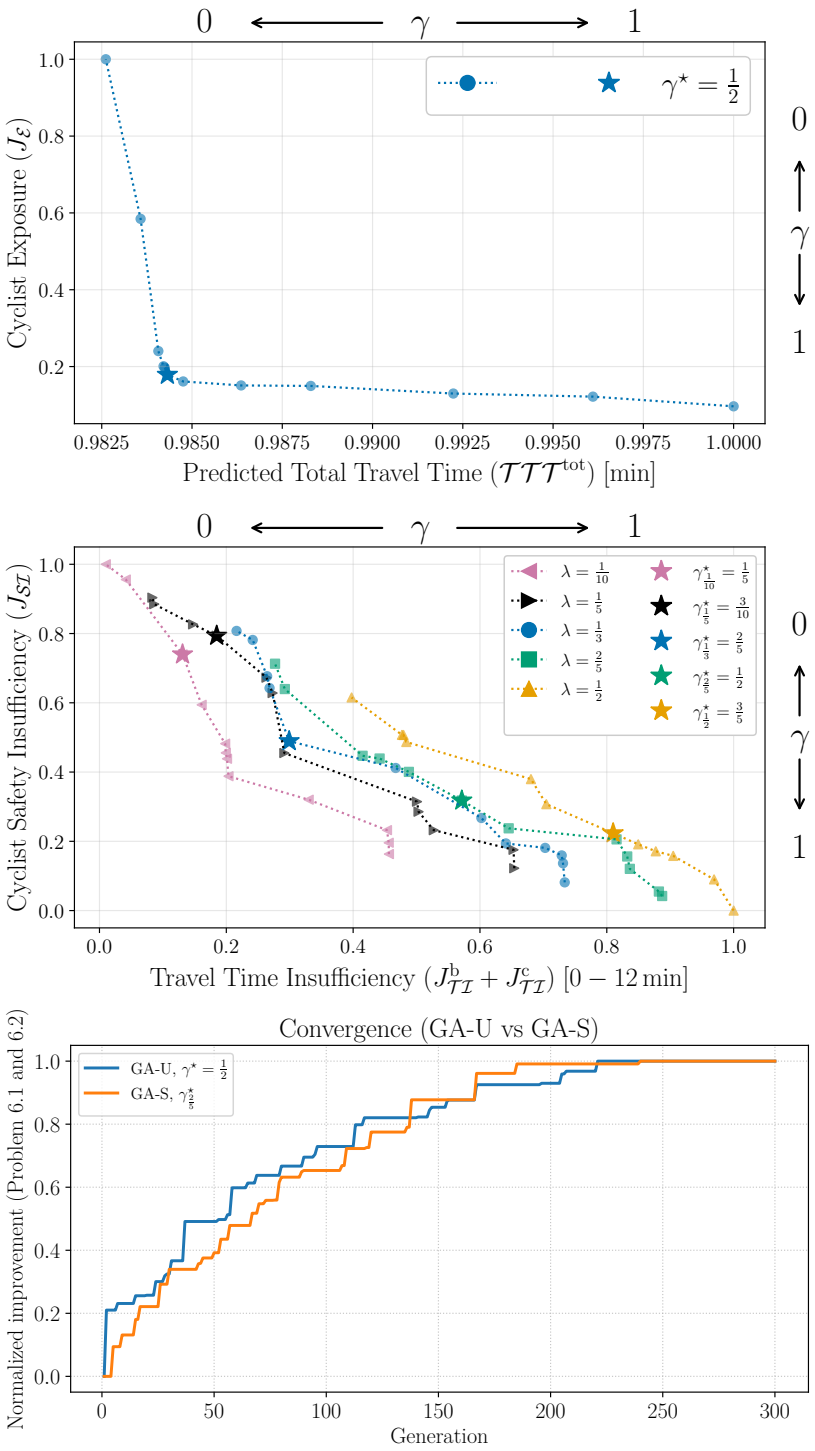

Figure 6.5: Normalized travel time versus cyclist exposure. The top panel shows the Pareto fronts obtained from the trade-off Problem 6.1 under varying values of $\gamma \in \Lambda$ and $B \approx 0.15 B_{\max}$. The medium panel presents the corresponding results for the sufficientarian formulation Problem 6.2. And the bottom row present the convergence analysis for the utilitarian optimal and sufficientarian for $\lambda = \frac{2}{5}$.

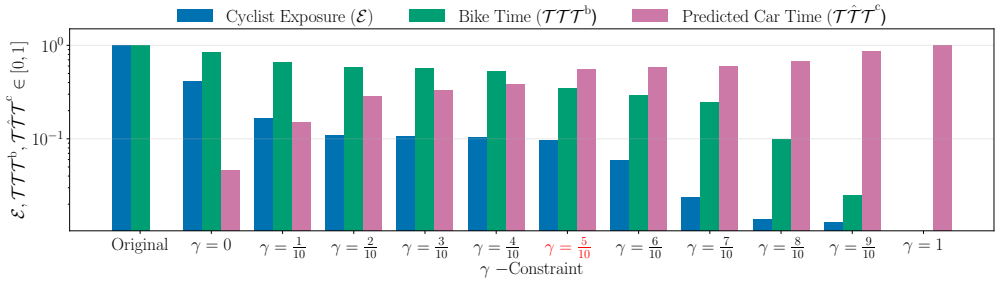

(a) Utilitarian perspective

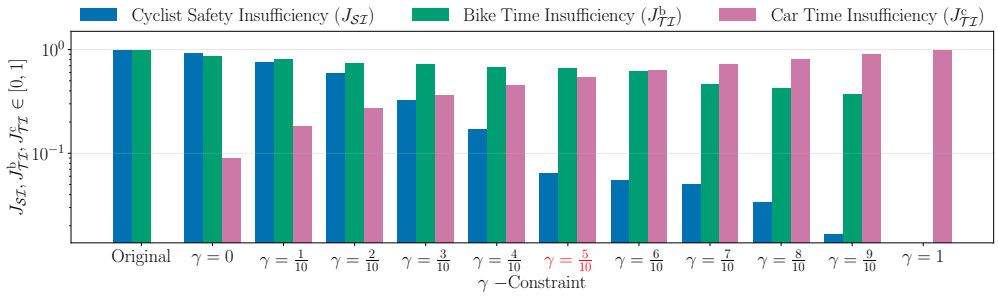

(b) Sufficientarian perspective $\lambda = \frac{2}{5}$

Figure 6.6: Normalized cyclist exposure, travel time, safety path insufficiency, cyclist and drivers trip insufficiency trade-off for $\lambda = \frac{2}{5}$.

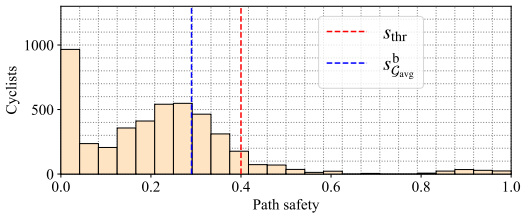

(a) Original

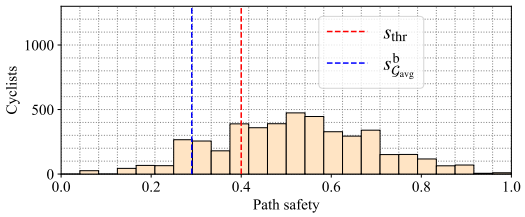

(b) GA-U

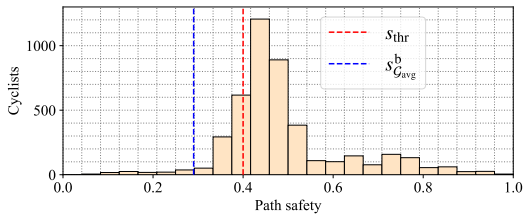

(c) GA-S

Figure 6.7: Path safety versus the number of cyclists above and below the threshold $\lambda = \frac{2}{5}$ under a budget of $B \approx 0.15 \cdot B_{\max}$. The dotted blue and red lines indicate the relative path-safety threshold s_{thr} and the average safety score of the entire network $S^{\mathrm{b}}_{\mathcal{G}_{\mathrm{avg}}}$, respectively.

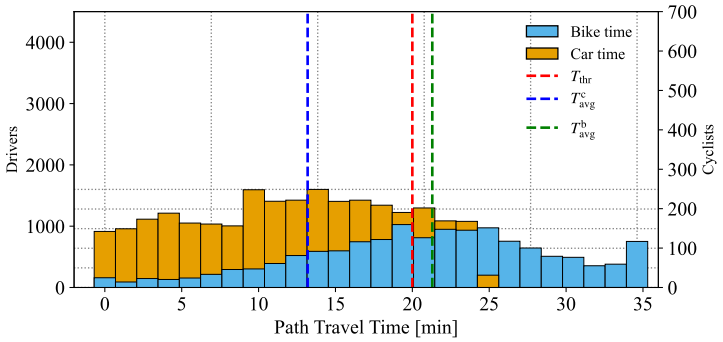

(a) Original

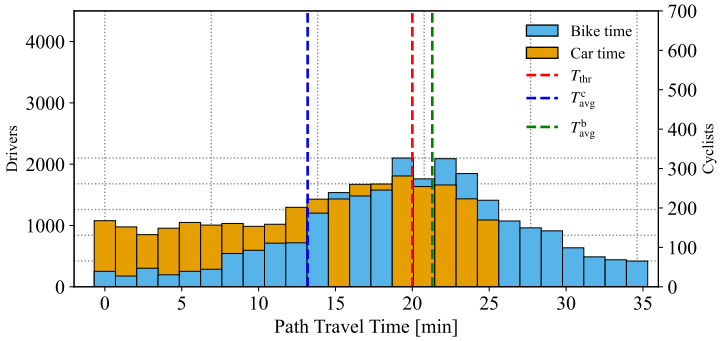

(b) GA-U

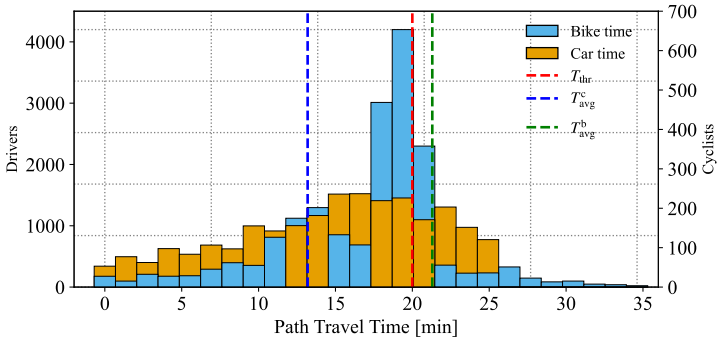

(c) GA-S

Figure 6.8: Drivers (left y-axis) and cyclists (right y-axis) trip insufficiency versus the number of individuals above and below the threshold $\lambda = \frac{2}{5}$ under a budget of $B^*_\lambda \approx 0.15 \cdot B_{\max}$ for the different optimization algorithms. The dotted red, blue, and green lines indicate the trip-insufficiency threshold T_{thr}, the average bicycle trip insufficiency $T^{\text{b}}_{\text{avg}}$, and the average car trip insufficiency $T^{\text{c}}_{\text{avg}}$ of the entire network, respectively.

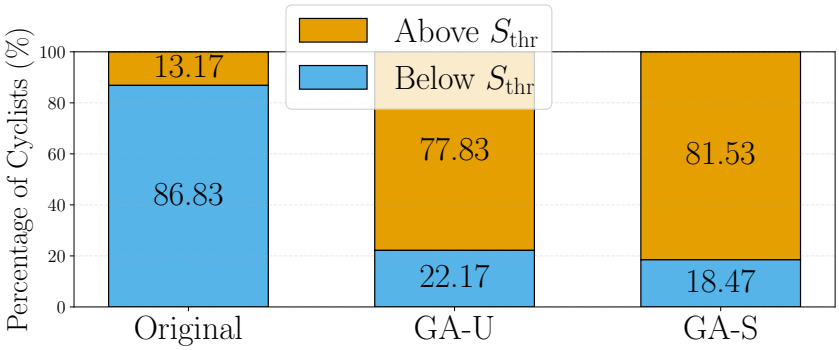

(a) Path-safety insufficiency $\lambda = \frac{2}{5}$

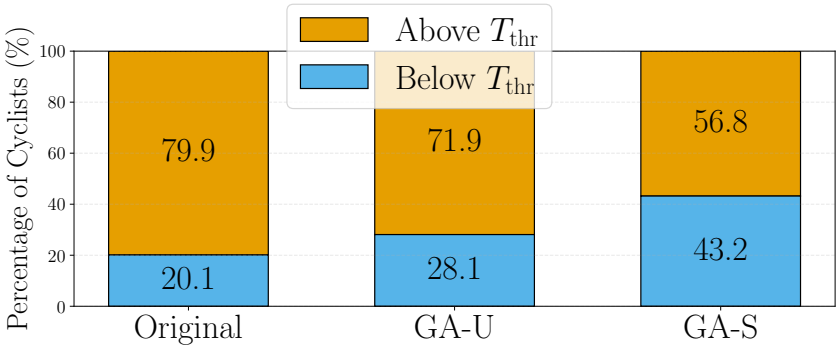

(b) Cyclist trip insufficiency $\lambda = \frac{2}{5}$

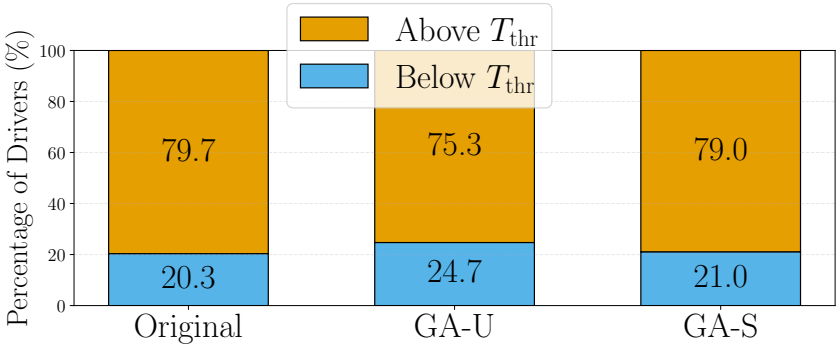

(c) Drivers trip insufficiency $\lambda = \frac{2}{5}$

Figure 6.9: Normalized path-safety insufficiency, bike path-time insufficiency and car path-time insufficiency for $\lambda = \frac{2}{5}$.

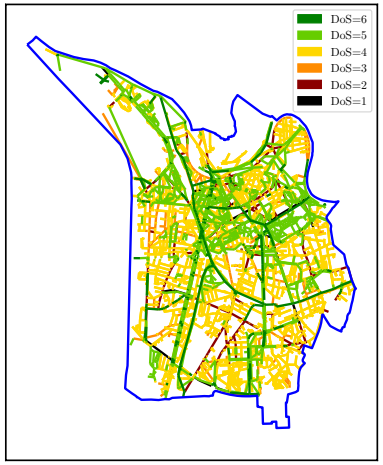

(a) Original DoS, $\overline{D}^{\mathrm{b}} = 4.41$, $P^{\mathrm{b}} = 16.49$ %

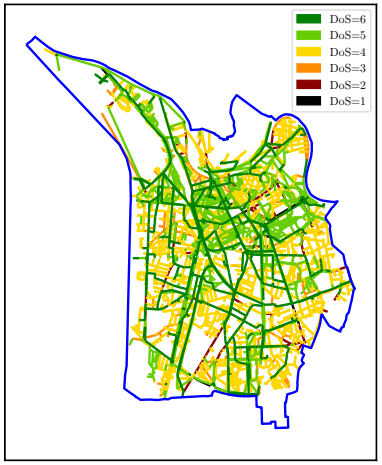

(b) MLCC DoS, $\overline{D}^{\mathrm{b}} = 4.29$ km, $P^{\mathrm{b}} = 20.07$ %.

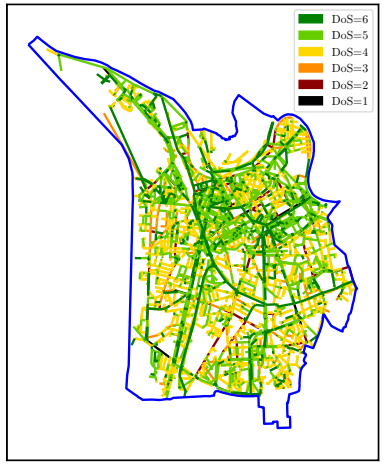

(c) GA-U DoS, $\overline{D}^{\mathrm{b}} = 4.36$, $P^{\mathrm{b}} = 18.05$ %.

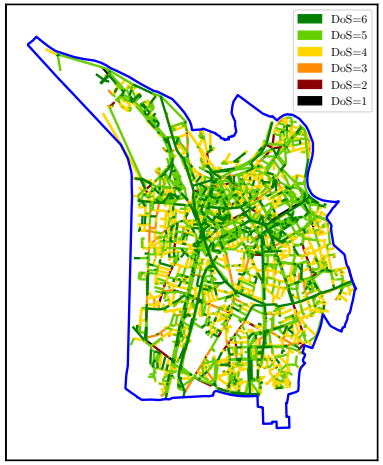

(d) GA-S DoS, $\overline{D}^{\mathrm{b}} = 4.31$, $P^{\mathrm{b}} = 19.48$ %.

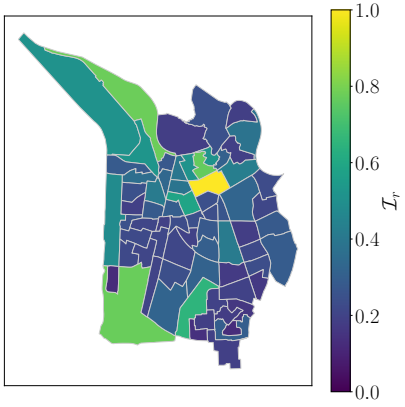

(e) Original safety path insufficiency

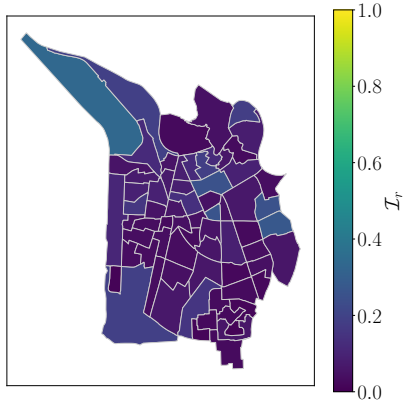

(f) MLCC safety path insufficiency

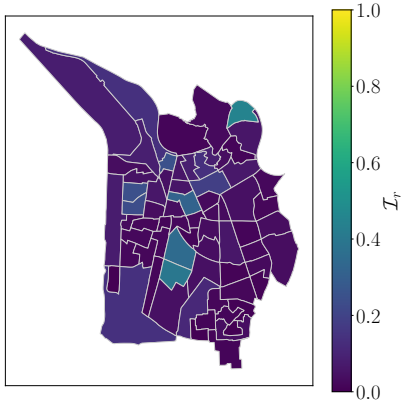

(g) GA-U safety path insufficiency

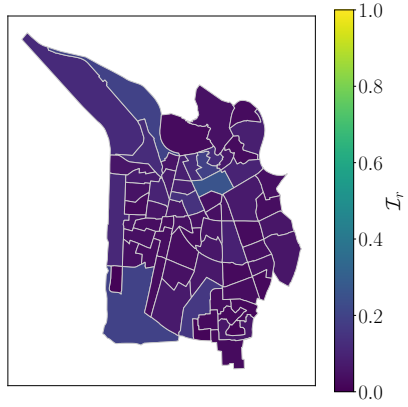

(h) GA-S safety path insufficiency

Figure 6.10: DoS configurations (top row) and regional safety path insufficiency (second row) for the original network and the MLCC, GA-U, and GA-S upgrade plans. All methods use $\lambda = \frac{2}{5}$ and budget $B \approx 0.15\, B_{\max}$.

(a) Original bike travel time insufficiency

(b) MLCC bike travel time insufficiency

(c) GA-U bike travel time insufficiency

(d) GA-S bike travel time insufficiency

(e) Original car travel time insufficiency

(f) MLCC car travel time insufficiency

(g) GA-U car travel time insufficiency

(h) GA-S car travel time insufficiency

Figure 6.11: Regional bicycle travel time insufficiency (top row) and car travel time insufficiency (bottom row) for the original network and the MLCC, GA-U, and GA-S upgrade plans. All methods use $\lambda = \frac{2}{5}$ and budget $B \approx 0.15\, B_{\max}$.

To illustrate these effects, Fig. 6.7 shows the relationship between path-safety scores and the number of cyclists for each method. Both MLCC and GA-S lead to a larger share of cyclists exceeding the safety threshold. Fig. 6.8 extends this analysis to both cyclists and drivers, showing the number of individuals above and below the time-accessibility threshold. While MLCC and GA-U effectively reduce cyclist accessibility unfairness, only GA-S allows more car drivers to remain below the time threshold, minimizing travel-time excess despite the reduced car-lane capacity caused by bike-lane upgrades. Fig. 6.9 illustrates, for $\lambda = \frac{2}{5}$, the percentage of cyclists above and below the safety threshold. Both MLCC and GA-S improve path safety and cyclist travel time; however, only GA-S mitigates the impact on car travel times, resulting in the lowest increase in travel time for motorized traffic. Finally, Fig. 6.10 and Fig. 6.11 compare the original DoS distribution and regional safety insufficiency before and after applying the MLCC, GA-U, and GA-S optimization strategies ($\lambda = 40\%$). MLCC concentrates the budget on connecting high-insufficiency regions to the safest network components, while GA-U distributes investments more evenly. Both GA-S and MLCC achieve the largest reductions in regional safety insufficiency and bicycle accessibility unfairness, but only GA-S limits the negative accessibility impacts on cars.

Fig. 6.12 presents the results of the sufficientarian approach in Problem 6.2, specifically to calculate the upper level $\mathbb{U}_{\text{thr}}$ for the different candidate solutions and threshold values given by λ. We observe that when λ is small, safety insufficiency remains low. However, as λ increases, insufficiency also rises. Despite this, MLCC and GA-S repeatedly achieve the greatest reductions in safety insufficiency.

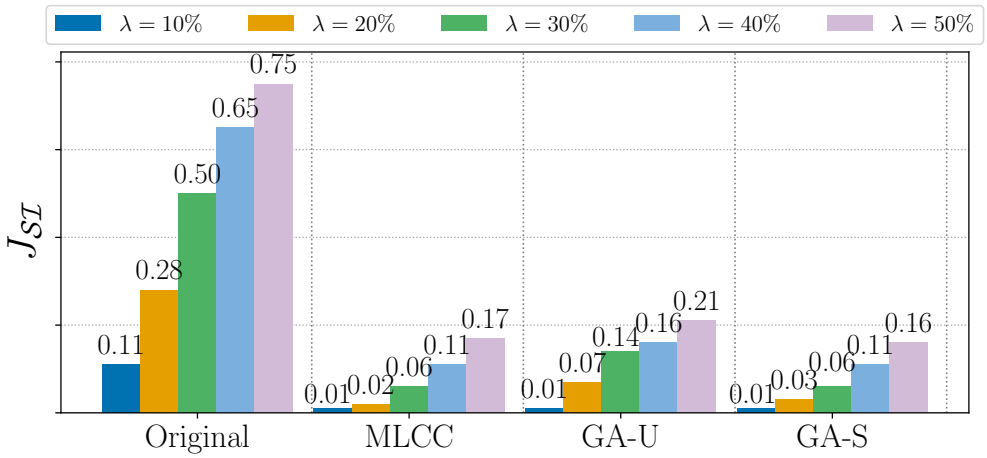

Figure 6.12: Safety insufficiency comparison of upper level $\mathbb{U}_{\text{thr}}$ of Problem 6.2, and the color legend indicates the λ threshold.

To assess whether the proposed framework also improves the overall efficiency of the multimodal system, we evaluate the total system travel time already defined in Eq. (6.4). Fig. 6.14 presents the total system travel time $\mathcal{TTT}^{\text{tot}}(\tilde{\boldsymbol{w}}^{\text{b}})$ across all configurations. In every optimization scenario, the $\mathcal{TTT}^{\text{tot}}$ is reduced compared to the baseline, primarily due to the modal shift from cars to bicycles. This reduction reflects that cyclists reach their destinations more efficiently within safer and more comfortable networks. Among the evaluated methods, the

sufficientarian approach achieves the greatest decrease in demand weighted total system travel time, indicating that incorporating a fairness-oriented perspective can simultaneously improve overall safety and reduce unfairness. Finally, Fig. 6.15 illustrates representative cyclist and driver paths for the different configurations, showing that only the sufficientarian framework enables both cyclists and drivers to remain below the time threshold. These results suggest that a carefully designed sufficientarian strategy can enhance not only fairness but also the overall temporal performance of the transport system, here represented by the two main travel modes, drivers and cyclists. This completely aligns with studies show that car drivers who correct their inaccurate knowledge of travel time with public transport and bike are more likely to use it in the future or at least consider in their choice set [137–139]. therefore, it is important to understand that using low-carbon modes can sometimes have comparable travel time to driving car, i.e., there is a potential for time-relevant low-carbon transportation.

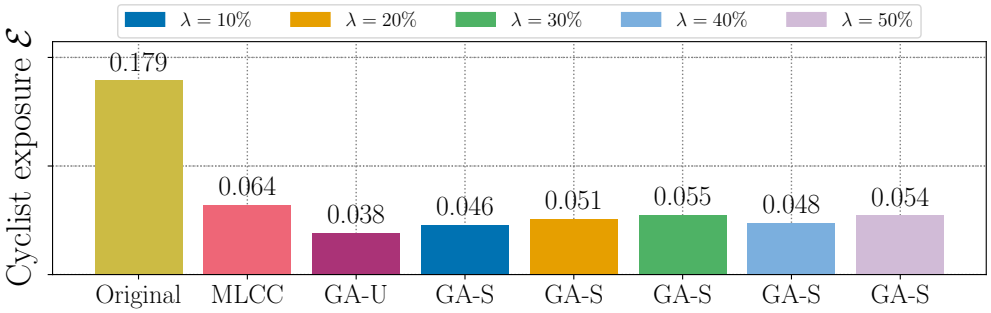

Figure 6.13: Normalized cyclist exposure.

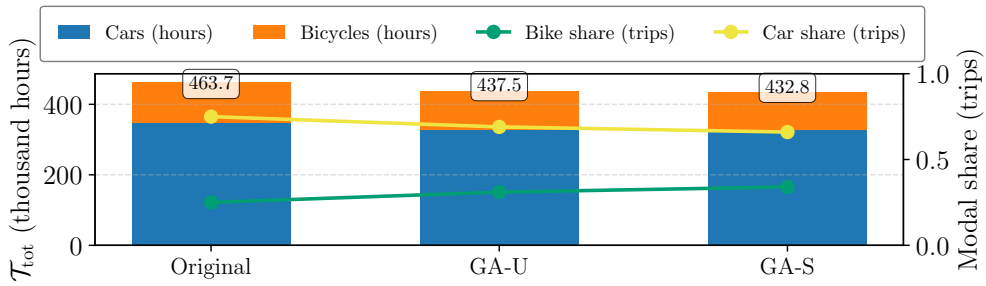

Figure 6.14: Total system travel time ($\mathcal{TTT}^{\text{tot}}$).

6.4 Concluding Remarks

This chapter introduced a multi-modal bi-level optimization framework for cycling-infrastructure planning that jointly accounts for (i) network upgrades through the DoS concept, (ii) induced modal shift from cars to bicycles, (iii) multi-modal traffic assignment with car user equilibrium responses, and (iv) transport-justice objectives. Applied to the Grenoble network, the framework demonstrated that substantial improvements in cycling safety can

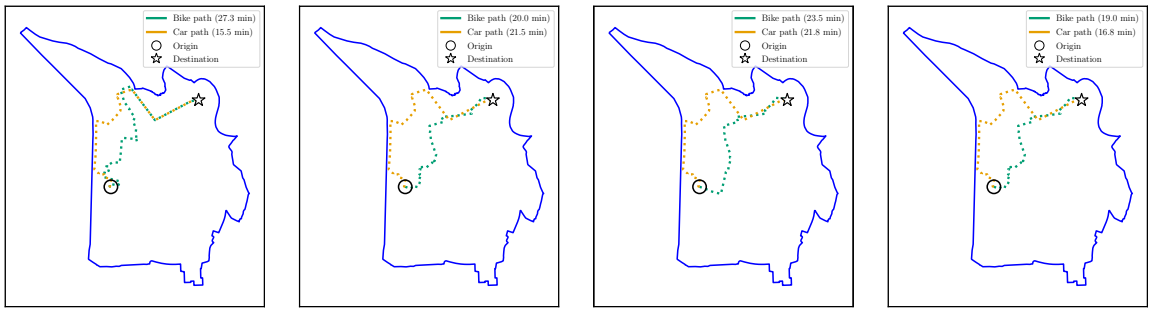

(a) Original DoS **(b)** MLCC DoS **(c)** GA-U DoS **(d)** GA-S DoS

Figure 6.15: Grenoble configuration paths under different optimizations algorithms.

be achieved while keeping the impacts on motor-vehicle performance limited, especially once demand feedback and post-upgrade traffic re-equilibration are considered. Beyond aggregate efficiency, the experiments showed that explicitly embedding justice criteria changes which interventions are prioritized: instead of concentrating gains where they are cheapest, the model can target gaps in minimum safety and accessibility. Among the tested strategies, the sufficientarian approach (GA-S) delivered the most balanced outcomes, simultaneously reducing cyclist and driver "insufficiency" shares while maintaining competitive system-level performance, thereby illustrating that equity-oriented planning can be compatible with operational feasibility.

Several extensions can strengthen the policy realism of the framework. First, we plan to enrich the mode-choice component with historical and contextual determinants of modal shift capturing how past infrastructure trajectories, local norms, policy packages, and investment intensity shape cycling uptake and car dependence. Incorporating such context-sensitive variables should improve the credibility and transferability of predicted shifts, in line with the broader view that mobility transitions are negotiated and path-dependent processes rather than purely instantaneous responses to travel-time changes [131].

Second, while on this chapter we report an efficiency proxy through the total system travel time $\mathcal{TTT}^{\text{tot}}$ (Fig. 6.14), a more complete efficiency baseline in the sense of transport appraisal requires a welfare or money-metric objective. Such a baseline would make it possible to quantify the cost of equity: for instance, how much welfare (or net present value) is foregone to achieve a given reduction in safety or accessibility shortfalls. We therefore plan to extend the objective beyond engineering indicators by integrating a benefit–cost perspective that values travel-time changes, safety improvements, and health gains against investment and external costs, enabling direct comparison between utilitarian and sufficientarian designs within a common accounting framework (e.g., NPV), in the spirit of Paulsen and Rich [61].

Finally, we will validate the full framework across several French cities to assess robustness to differences in urban form, network structure, and baseline cycling cultures, and to better understand how utilitarian and sufficientarian priorities translate into actionable, context-dependent investment strategies.

Chapter 7

Clustering and Comparing Safety in Cycling Networks

Contents

7.1 Introduction . . . 131
7.1.1 Summary of Results . . . 132
7.1.2 Related Work . . . 133
7.2 Model Design . . . 134
7.2.1 Cluster Network Definition . . . 135
7.2.2 Clustered Cyclist Distance . . . 135
7.2.3 Clustered Demand, Flow Assignment, and Exposure . . . 137
7.3 Clustering (Doc-CS) . . . 138
7.3.1 Mathematical Guarantees . . . 142
7.4 Network Comparison . . . 143
7.5 Benchmark Instances and Case Study . . . 144
7.6 Concluding Remarks . . . 149

7.1 Introduction

This final chapter is more technical in nature and addresses a methodological issue that becomes important when cycling networks are studied across different cities: *how can we compare safety-related properties of networks that differ in size, spatial coverage, and mapping detail?* Cities vary not only in the amount of infrastructure they have, but also in how their networks are represented in available data. As a result, direct comparisons based on the raw graphs can be hard to interpret, because part of the observed differences may come from differences in representation rather than from differences in safety conditions.

To support cross-city assessment in a clearer and more consistent way, we introduce a framework for fair and interpretable comparison of cycling networks. The main idea is to build a reduced representation of each city network that keeps the safety-relevant structure—such as the connectivity of highly protected links and exposure-related attributes—while limiting sensitivity to network size and data resolution. This provides a practical basis for benchmarking

cycling networks across cities, comparing alternative designs, and monitoring improvements over time using comparable indicators.

To address this need, we propose a mathematical framework that cluster large cycling networks into smaller, interpretable clustered-networks while preserving the safety- and flow-relevant information used throughout this thesis. The core idea is to replace a detailed road graph with a cluster-based representation obtained from a constrained, safety-aware k-means applied to a spectral embedding of the network. Each cluster becomes a clustered-node, and clustered-arcs summarize the connections between clusters. Crucially, the compression is not purely geometric or topological: it is designed to preserve the key attributes that determine cyclist route choice and exposure, including length, slope, degree of separation (DoS), traffic factors, bicycle flows, and exposure measures, while explicitly limiting information loss.

Building on these clustered networks, we then develop exposure-based kernels that encode the safety-critical structure of the network at the clustered-network level. These kernels allow us to compare cities through principled, label-invariant similarity measures, producing a notion of "distance" between urban cycling networks that is meaningful for safety analysis and robust to differences in network resolution. Experiments on real networks show that the clustered representations retain enough safety-relevant structure to support consistent and interpretable cross-city comparisons.

7.1.1 Summary of Results

The results show that the proposed clustering procedure can compress large cycling networks into fixed-size clustered-networks while retaining the safety-relevant structure needed for comparison. In particular, the constrained, safety-aware k-means on the spectral embedding produces clusters that remain consistent with the underlying network geometry and connectivity, so that clustered-nodes and clustered-arcs still represent meaningful corridors and inter-area links rather than purely geometric partitions. This makes the reduced graphs interpretable: the main "safe backbone" of each city (i.e., how highly protected segments connect and where they fragment) remains visible after reduction.

At the attribute level, experiments indicate that the clustering preserves the quantities that matter for cyclist experience and for the exposure model used throughout the thesis. Aggregated link properties—length, slope, DoS composition, traffic factors, cyclist flows, and exposure—remain aligned between the original network and its clustered representation, meaning that the clustered-network does not simply keep topology but also maintains the weighted safety information. Practically, this implies that network-wide summaries computed on the reduced graph (e.g., exposure concentration, connectivity of high-DoS infrastructure, and flow–exposure coupling) remain representative of the original city network.

Using these reduced networks, the exposure-based kernels let us compare cities in a way that is not dominated by network size or by how detailed the map data is. The similarity score does not depend on node labels, and it accounts for both the network structure and the safety-related weights (DoS, flow, exposure). This means that two cities are judged as

similar when they organize safe, low-exposure routes in a similar way—not simply when one city has a more detailed graph. In our experiments, the resulting distances are stable and easy to interpret: cities with comparable "safe backbones" end up close to each other, while cities with different exposure patterns are clearly separated, which is exactly what we need for benchmarking.

Overall, the chapter shows that combining clustering with exposure-based comparison gives a useful tool for planners. It puts different cities on the same scale, keeps the key safety information, and makes it possible to compare networks and track changes over time in a consistent and interpretable way.

7.1.2 Related Work

Our contribution connects to two main research streams: (i) clustering and community detection in transportation networks, and (ii) graph comparison.

Clustering and community detection: Cluster analysis and community detection are standard tools to reduce complexity and reveal structure in large datasets and networks. Classical k-means partitions points into k groups by minimizing within-cluster variance [140], while network-specific community detection methods identify groups of nodes that are more densely connected internally than externally. Representative approaches include divisive methods based on arc betweenness [141], greedy modularity maximization [142], spectral modularity-based algorithms [143], and physics-inspired techniques such as voltage-drop analogies [144].

In transportation, these ideas have been used to analyze traffic regimes, detect disruptions, or partition networks into control regions. For example, k-means has been applied to identify traffic-flow phases [145] and disruptions [146], while clustering has also supported behavioral characterization from mobility traces [147]. For spatially embedded networks, community detection can support network management and equity analyze; e.g., [148] refine Info map-based partitions to satisfy contiguity and compactness constraints. In cycling-related work, demand-based communities have been used to guide bike network design [102], and data-driven strategies have been proposed to connect fragmented bicycle networks under budget constraints [63].

Compared with these studies, our clustering procedure is explicitly safety-aware: it is designed to compress cycling networks while preserving cyclist-oriented attributes that directly shape perceived safety and comfort, such as exposure, DoS, slope, traffic conditions, and flow. To the best of our knowledge, existing clustering and community-detection approaches are not tailored to preserve these multi-attribute safety indicators when reducing large cycling graphs for cross-city comparison.

Network comparison: Comparing transport networks can be viewed as a relaxed version of graph isomorphism: rather than testing whether two graphs are identical up to relabeling, we quantify how different their structures are under a chosen distance [149]. Simple approaches compare adjacency matrices via standard norms, but these are highly sensitive to node labeling

and local perturbations [150]. More robust strategies rely on structural descriptors, such as global statistics (diameter, clustering, average path length), or role-based features [151, 152]. Spectral signatures based on adjacency or Laplacian eigenvalues provide label-invariant representations and often correlate with edit-based notions of graph difference [153]. Other approaches use graph limits and cut distances to compare networks in a size-agnostic manner [154].

A complementary line of work embeds graphs into reproducing kernel Hilbert spaces via graph kernels, enabling principled similarity measures between structured objects. Random-walk and marginalized kernels are particularly attractive because they provide tractable feature spaces without requiring explicit subgraph enumeration [155, 156]. These ideas connect naturally to optimal transport, where Wasserstein distances offer geometrically meaningful metrics between distributions and have become central tools in modern applied mathematics [157–159].

In our setting, network comparison must satisfy two requirements that are not simultaneously addressed by existing methods: it must be robust to differences in network size/resolution, and it must remain sensitive to cycling-safety attributes (DoS, exposure, slope, and flow) rather than relying only on unweighted topology. Our framework addresses this gap by (i) clustering cycling networks under explicit safety constraints into comparable clustered-networks, and (ii) comparing these reduced representations through exposure-based kernels that encode both structure and safety-relevant weights. To the best of our knowledge, there is currently no operational framework that combines these two components to support consistent, safety-oriented benchmarking of cycling networks across cities.

7.2 Model Design

This section recalls the notation and core definitions used to model the cycling network. In particular, we introduce an exposure indicator that will later serve as the basis for cross-network comparison.

To obtain a safety measure that remains meaningful after graph compression and is not overly sensitive to network size or mapping resolution, we work with an *average* exposure rather than a total exposure. Total exposure naturally increases with the number of arcs and with the level of geometric detail in the data (e.g., when a single street is represented by many short OpenStreetMap segments), which can bias comparisons and distort the clustering objective. Averaging over arcs yields a normalized quantity that better reflects the typical exposure level experienced in the network and reduces sensitivity to representation artifacts. We therefore define the average cyclist exposure of the original network as

$$\mathcal{E}^{\text{orig}} = \frac{1}{|\mathcal{A}^{\text{b}}|} \sum_{a \in \mathcal{A}^{\text{b}}} \phi_a^{\text{b}} \cdot f_a^{\text{b}} \cdot d_a^{\text{b}}. \tag{7.1}$$

We also denote the aggregate cyclist flow in $\mathcal{G}^{\mathrm{b}}$ by

$$\Phi^{\mathrm{orig}} \;=\; \sum_{a\in\mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}}. \tag{7.2}$$

7.2.1 Cluster Network Definition

We now present the clustering framework used to construct the reduced (clustered-network) representation of the cycling graph.

Let the set $\{C_1,\dots,C_k\}$ be a partition of the node set $\mathcal{V}^{\mathrm{b}}$ into k non-empty clusters. We define the clustered network as $\tilde{\mathcal{G}}^{\mathrm{b}} = (\tilde{\mathcal{V}}^{\mathrm{b}}, \tilde{\mathcal{A}}^{\mathrm{b}}, \tilde{\mathcal{W}}^{\mathrm{b}})$, where the clustered-node set is $\tilde{\mathcal{V}}^{\mathrm{b}} = \{C_1,\dots,C_k\}$. For any pair of clusters $(C_r, C_{r'})$, we define the set of arcs between them as

$$\mathcal{A}^{\mathrm{b}}(C_r, C_{r'}) \;=\; \big\{(u,v)\in\mathcal{A}^{\mathrm{b}} \;\big|\; u\in C_r,\; v\in C_{r'}\big\}.$$

We include a clustered-arc $\tilde{a} = (C_r, C_{r'})$ in $\tilde{\mathcal{A}}^{\mathrm{b}}$ if and only if $\mathcal{A}^{\mathrm{b}}(C_r, C_{r'}) \neq \varnothing$. The clustered arc weights are split into *inter*- and *intra*-cluster components. For $r \neq r'$, the inter-cluster DoS weight between C_r and $C_{r'}$ is defined as

$$\tilde{w}^{\mathrm{inter}}_{(C_r,C_{r'})} \;=\; \min_{a\in\mathcal{A}^{\mathrm{b}}(C_r,C_{r'})} w_a^{\mathrm{b}}, \tag{7.3}$$

i.e., the minimum DoS (highest-risk) connection between the two clusters. For intra-cluster connections, we write

$$\mathcal{A}^{\mathrm{b}}(C_r) \;:=\; \mathcal{A}^{\mathrm{b}}(C_r, C_r) \;=\; \big\{(u,v)\in\mathcal{A}^{\mathrm{b}} \;\big|\; u\in C_r,\; v\in C_r\big\},$$

and define the intra-cluster DoS weight as

$$\tilde{w}^{\mathrm{intra}}_{C_r} \;=\; \min_{a\in\mathcal{A}^{\mathrm{b}}(C_r)} w_a^{\mathrm{b}}, \tag{7.4}$$

which represents the worst-case (highest-risk) connection within cluster C_r.

To enable flow and exposure analysis on the clustered network $\tilde{\mathcal{G}}^{\mathrm{b}}$, we need to redefine the AoN traffic assignment approach for this aggregated representation. This requires establishing new rules for path selection and flow distribution across the clustered structure.

7.2.2 Clustered Cyclist Distance

Let $\tilde{p}_{(C_r,C_{r'})} = (\tilde{a}_1,\dots,\tilde{a}_m)$ denote a shortest path between clusters C_r and $C_{r'}$ in the clustered network $\tilde{\mathcal{G}}^{\mathrm{b}}$, where each clustered-arc $\tilde{a}_\ell \in \tilde{p}_{(C_r,C_{r'})}$ connects two clusters. We write $\mathcal{C}(C_r, C_{r'}) \;\subseteq\; \tilde{\mathcal{V}}^{\mathrm{b}}$ for the set of clusters visited along this path (including C_r and $C_{r'}$). The clustered cyclist distance between C_r and $C_{r'}$ is defined as

$$\tilde{b}_{(C_r,C_{r'})} \;=\; \underbrace{\sum_{\tilde{a}\in\tilde{p}_{(C_r,C_{r'})}} \tilde{b}^{\mathrm{inter}}_{\tilde{a}}}_{\text{inter-cluster}} \;+\; \underbrace{\sum_{C_q\in\mathcal{C}(C_r,C_{r'})} \tilde{b}^{\mathrm{intra}}_{C_q}}_{\text{intra-cluster}}. \tag{7.5}$$

For each clustered-arc $\tilde{a} = (C_r, C_{r'}) \in \tilde{\mathcal{A}}^{\mathrm{b}}$ with $r \neq r'$, the inter-cluster contribution is

$$\tilde{b}^{\mathrm{inter}}_{\tilde{a}} \;=\; \big(\tilde{d}^{\mathrm{inter}}_{\tilde{a}} \cdot \tilde{s}^{\mathrm{inter}}_{\tilde{a}} \cdot (\tilde{f}^{\mathrm{inter}}_{\tilde{a}})^{\rho}\big) \;+\; \tilde{\measuredangle}^{\mathrm{inter}}_{\tilde{a}} \cdot \kappa, \tag{7.6}$$

where the aggregated distance, slope, traffic factor and angular change are defined from the underlying arcs between C_r and $C_{r'}$:

$$\tilde{d}^{\mathrm{inter}}_{\tilde{a}} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} d^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} \phi^{\mathrm{b}}_a}, \quad \text{(distance)} \tag{7.7}$$

$$\tilde{s}^{\mathrm{inter}}_{\tilde{a}} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} s^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} \phi^{\mathrm{b}}_a}, \quad \text{(slope)} \tag{7.8}$$

$$\tilde{f}^{\mathrm{inter}}_{\tilde{a}} = f_a\big(\tilde{w}^{\mathrm{inter}}_{(C_r, C_{r'})}\big), \quad \text{(traffic factor)} \tag{7.9}$$

$$\tilde{\measuredangle}^{\mathrm{inter}}_{\tilde{a}} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} \measuredangle^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r, C_{r'})} \phi^{\mathrm{b}}_a}, \quad \text{(angular change).} \tag{7.10}$$

For each cluster C_r, we define the set of intra-cluster arcs

$$\mathcal{A}^{\mathrm{b}}(C_r) \;:=\; \mathcal{A}^{\mathrm{b}}(C_r, C_r) \;=\; \big\{(u,v) \in \mathcal{A}^{\mathrm{b}} \,\big|\, u \in C_r,\; v \in C_r\big\}.$$

The intra-cluster contribution of C_r is

$$\tilde{b}^{\mathrm{intra}}_{C_r} \;=\; \big(\tilde{d}^{\mathrm{intra}}_{C_r} \cdot \tilde{s}^{\mathrm{intra}}_{C_r} \cdot (\tilde{f}^{\mathrm{intra}}_{C_r})^{\rho}\big) \;+\; \tilde{\measuredangle}^{\mathrm{intra}}_{C_r} \cdot \kappa, \tag{7.11}$$

with

$$\tilde{d}^{\mathrm{intra}}_{C_r} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} d^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} \phi^{\mathrm{b}}_a}, \quad \text{(distance)} \tag{7.12}$$

$$\tilde{s}^{\mathrm{intra}}_{C_r} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} s^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} \phi^{\mathrm{b}}_a}, \quad \text{(slope)} \tag{7.13}$$

$$\tilde{f}^{\mathrm{intra}}_{C_r} = f_a\big(\tilde{w}^{\mathrm{intra}}_{C_r}\big), \quad \text{(traffic factor)} \tag{7.14}$$

$$\tilde{\measuredangle}^{\mathrm{intra}}_{C_r} = \frac{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} \measuredangle^{\mathrm{b}}_a \cdot \phi^{\mathrm{b}}_a}{\sum_{a \in \mathcal{A}^{\mathrm{b}}(C_r)} \phi^{\mathrm{b}}_a}, \quad \text{(angular change).} \tag{7.15}$$

Remark 7.1: Remark

Our clustering methodology and cyclist-distance formulation are designed to capture essential aspects of cycling behaviour while remaining computationally tractable. Key design choices include:

- Using minimum DoS values for inter- and intra-cluster connections to follow a conservative, bottleneck-preserving perspective: a single high-risk segment can dominate perceived safety along a route.

- Aggregating distance, slope, traffic factor, and angular change via flow-weighted averages, so that frequently used segments dominate cluster-level characteristics and better reflect typical cyclist experience.

7.2.3 Clustered Demand, Flow Assignment, and Exposure

Given the original graph $\mathcal{G}^{\mathrm{b}}$ and its clustered representation $\tilde{\mathcal{G}}^{\mathrm{b}}$. The clustered demand between two clusters $C_r, C_{r'} \in \mathcal{C}$ is obtained by aggregating the original OD demands $q^{\mathrm{b}}_{(u,v)}$ with origins in C_r and destinations in $C_{r'}$:

$$\tilde{q}_{(C_r,C_{r'})} := \sum_{\substack{u\in C_r \\ v\in C_{r'}}} q^{\mathrm{b}}_{(u,v)} \cdot \mathbb{I}_{(u,v)\in\mathcal{A}^{\mathrm{b}}}, \tag{7.16}$$

where $\mathbb{I}_{(u,v)\in\mathcal{A}^{\mathrm{b}}}$ is an indicator function that ensures only valid arcs in the original graph are considered. For the AoN assignment on the clustered graph, let $\tilde{p}_{(C_r,C_{r'})}$ denote a shortest path between C_r and $C_{r'}$ in $\tilde{\mathcal{G}}^{\mathrm{b}}$, computed with respect the clustered cyclist distance. We define the inter-cluster flow on a clustered-arc $\tilde{a} \in \tilde{\mathcal{A}}^{\mathrm{b}}$ as

$$\tilde{\phi}^{\mathrm{inter}}_{\tilde{a}} := \sum_{(r,r')} \tilde{q}_{(C_r,C_{r'})} \cdot \mathbb{I}_{\tilde{a}\in\tilde{p}_{(C_r,C_{r'})}}, \tag{7.17}$$

where $\mathbb{I}_{\tilde{a}\in\tilde{p}_{(C_r,C_{r'})}}$ indicates whether the clustered-arc $\tilde{a}$ lies on the shortest path between C_r and $C_{r'}$. Similarly, we define the intra-cluster flow associated with a cluster $C_r \in \mathcal{C}$ as the total clustered demand whose shortest path traverses C_r:

$$\tilde{\phi}^{\mathrm{intra}}_{C_r} := \sum_{(r_1,r_2)} \tilde{q}_{(C_{r_1},C_{r_2})} \cdot \mathbb{I}_{C_r\in\tilde{p}_{(C_{r_1},C_{r_2})}}, \tag{7.18}$$

where $\mathbb{I}_{C_r\in\tilde{p}_{(C_{r_1},C_{r_2})}}$ is 1 if the path between C_{r_1} and C_{r_2} passes through cluster C_r, and 0 otherwise. With all these, finally, on the clustered network $\tilde{\mathcal{G}}^{\mathrm{b}}$, we define the total cyclist exposure as the average of the inter- and intra-cluster contributions:

$$\begin{aligned} \mathcal{E}^{\mathrm{clust}} &:= \frac{1}{|\tilde{\mathcal{A}}^{\mathrm{b}}| + |\mathcal{C}|}\Big(\mathcal{E}^{\mathrm{inter}} + \mathcal{E}^{\mathrm{intra}}\Big), \\ \Phi^{\mathrm{clust}} &:= \Phi^{\mathrm{inter}} + \Phi^{\mathrm{intra}}. \end{aligned} \tag{7.19}$$

Here, the inter- and intra-cluster exposure terms are

$$\begin{aligned} \mathcal{E}^{\mathrm{inter}} &:= \sum_{\tilde{a}\in\tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\phi}^{\mathrm{inter}}_{\tilde{a}} \cdot f^{\mathrm{inter}}_{\tilde{a}} \cdot d^{\mathrm{inter}}_{\tilde{a}}, \\ \mathcal{E}^{\mathrm{intra}} &:= \sum_{r=1}^{k} \tilde{\phi}^{\mathrm{intra}}_{C_r} \cdot f^{\mathrm{intra}}_{C_r} \cdot d^{\mathrm{intra}}_{C_r}, \end{aligned} \tag{7.20}$$

and the corresponding total flows are

$$\begin{aligned}\Phi^{\text{inter}} &:= \sum_{\tilde{a}\in\tilde{\mathcal{A}}^{\text{b}}} \tilde{\phi}^{\text{inter}}_{\tilde{a}},\\ \Phi^{\text{intra}} &:= \sum_{r=1}^{k} \tilde{\phi}^{\text{intra}}_{C_r}.\end{aligned} \tag{7.21}$$

The quantities $d^{\text{inter}}_{\tilde{a}}$, $f^{\text{inter}}_{\tilde{a}}$ and $d^{\text{intra}}_{C_r}$, $f^{\text{intra}}_{C_r}$ are the clustered distance and traffic-factor terms defined in the previous subsection for inter- and intra-cluster connections, respectively.

7.3 Clustering (Doc-CS)

To enable fair comparison between two cycling networks of different sizes, we cluster the larger network so that its reduced representation has the same number of clustered-nodes as the smaller one. The construction of $\tilde{\mathcal{G}}^{\text{b}}$ is designed to preserve safety-relevant information including DoS, distance, slope, traffic exposure, and cyclist flows—so that the clustered network remains comparable to the original in safety terms. And the algorithm is called DoS-CS Algorithm.

We begin by defining a safety-aware symmetric weighted adjacency (similarity) matrix $A \in \mathbb{R}^{|\mathcal{V}^{\text{b}}|\times|\mathcal{V}^{\text{b}}|}$:

$$A_{uv} = \begin{cases} w^{\text{b}}_{uv} \cdot \exp\!\Big(-\dfrac{d^{\text{b}}_{uv}}{\gamma}\Big)\left(\dfrac{\phi^{\text{sym}}_{uv}}{\max_{(i,j)\in\mathcal{A}^{\text{b}}}\phi^{\text{sym}}_{ij}}\right)^{\alpha} & \text{if } \{u,v\}\in\mathcal{A}^{\text{b}},\\ 0 & \text{otherwise,}\end{cases} \tag{7.22}$$

where taking the elements of $\mathcal{A}^{\text{b}}$ as undirected arcs associated with the cycling network, w^{b}_{uv} is the DoS weight of the link between u and v, and d^{b}_{uv} is its spatial length. The term ϕ^{sym}_{uv} is a symmetric flow proxy obtained from the directed arc flows, e.g. $\phi^{\text{sym}}_{uv} := \frac{1}{2}(\phi^{\text{b}}_{(u,v)} + \phi^{\text{b}}_{(v,u)})$ (with $\phi^{\text{b}}_{(u,v)} = 0$ if $(u,v)\notin\mathcal{A}^{\text{b}}$). The parameter $\gamma > 0$ controls the spatial locality of the similarity, and $\alpha \geq 0$ regulates the influence of high-flow connections. From A, we form the normalized graph Laplacian

$$\mathcal{L} = I - D^{-1/2}AD^{-1/2}, \qquad D_{uu} = \sum_{v\in\mathcal{V}^{\text{b}}} A_{uv}, \tag{7.23}$$

and compute a spectral embedding from the leading eigenvectors of $\mathcal{L}$. Concretely, we solve the eigenvalue problem

$$\mathcal{L}\,u_i = \lambda_i u_i, \qquad 0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{|\mathcal{V}^{\text{b}}|},$$

where $u_i \in \mathbb{R}^{|\mathcal{V}^{\text{b}}|}$ are orthonormal eigenvectors. The intuition is that the smallest eigenvalues correspond to smooth eigenvectors on the graph, i.e., vectors that vary little across strongly connected node pairs (high A_{uv}). Hence, these eigenvectors provide low-dimensional coordinates that reveal community structure.

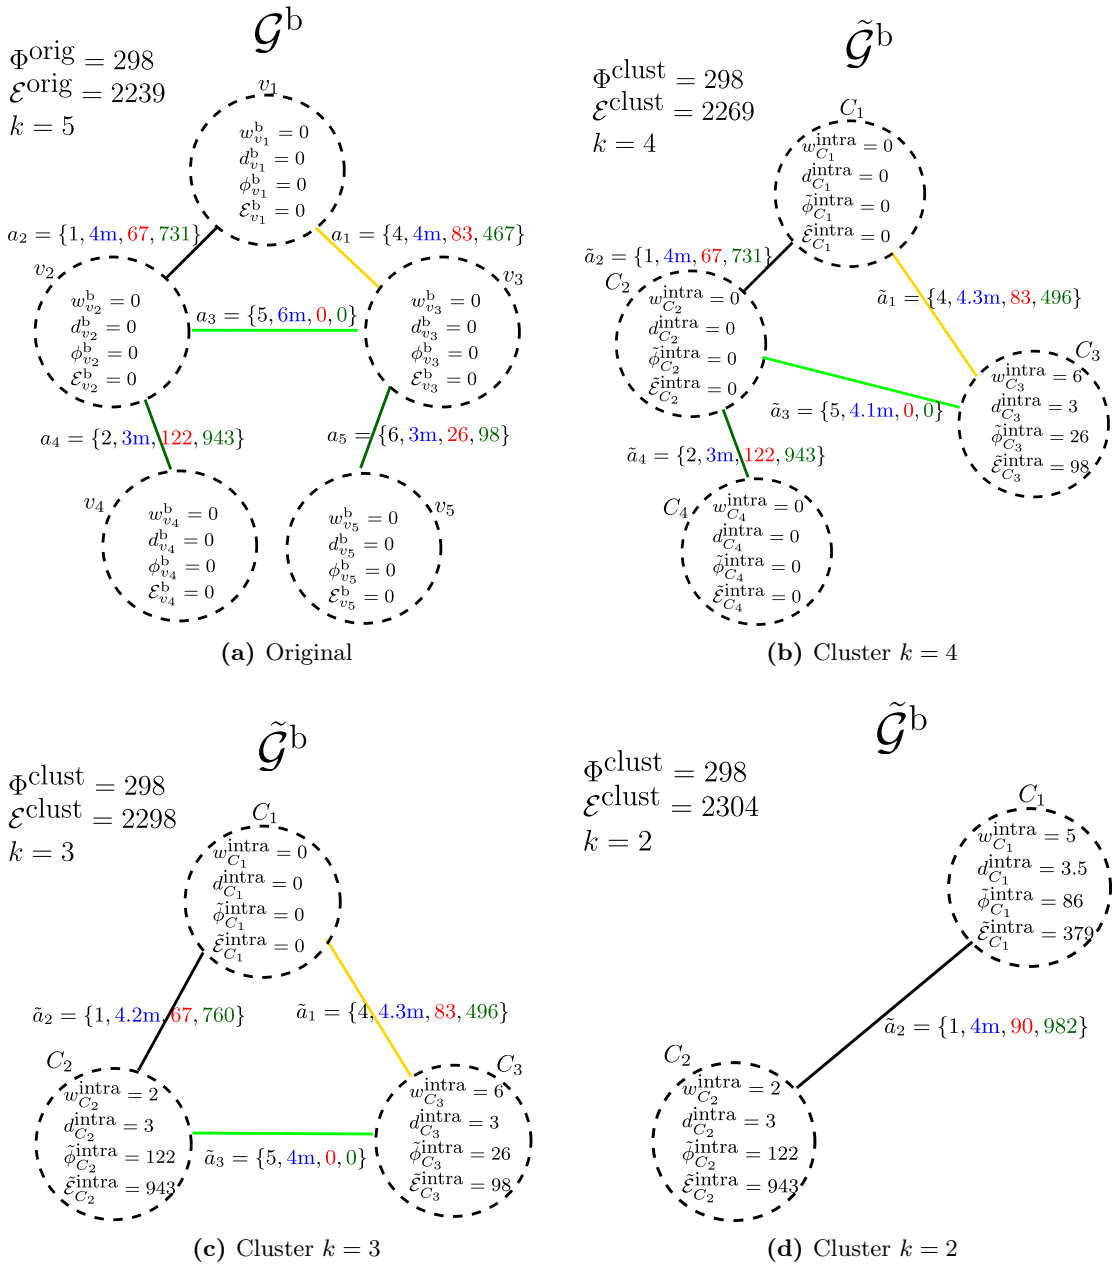

(a) Original **(b)** Cluster $k = 4$

(c) Cluster $k = 3$ **(d)** Cluster $k = 2$

Figure 7.1: DoS-CS Algorithm for toy example where $a = \{w_a, d_a(\mathrm{m}), \phi_a, \mathcal{E}_a\}$, $\tilde{a}^{\mathrm{inter}} = \left\{w^{\mathrm{inter}}_{\tilde{a}}, d^{\mathrm{inter}}_{\tilde{a}}(\mathrm{m}), \phi^{\mathrm{inter}}_{\tilde{a}}, \mathcal{E}^{\mathrm{inter}}_{\tilde{a}}\right\}$, and $C^{\mathrm{intra}} = \left\{d^{\mathrm{intra}}_{C}, d^{\mathrm{intra}}_{C}(\mathrm{m}), \phi^{\mathrm{inter}}_{C}, \mathcal{E}^{\mathrm{intra}}_{C}\right\}$

To embed each node $u \in \mathcal{V}^{\mathrm{b}}$ into $\mathbb{R}^m$, we stack the first m informative eigenvectors (typically excluding the trivial constant eigenvector u_1) into the matrix

$$U = [u_2, \ldots, u_{m+1}] \in \mathbb{R}^{|\mathcal{V}^{\mathrm{b}}| \times m},$$

and define the embedded point of node u as the u-th row of U (in our case $m = 2$):

$$x_u := U_{u,:} \in \mathbb{R}^m.$$

We then apply k-means in this embedding, with additional constraints to limit the distortion of flow and exposure in the clustered representation.

Let $\mathcal{E}^{\text{orig}}$ and Φ^{orig} denote the total cyclist flow and exposure on $\mathcal{G}^{\text{b}}$ (Eq. (7.1) and (7.2)), and let $\mathcal{E}^{\text{clust}}$ and Φ^{clust} be their clustered counterparts (defined in the previous subsection and Eq. (7.19)). These quantities are used to enforce that the clustering preserves the safety- and usage-relevant structure of the original network. The clustering is performed by a constrained k-means procedure (Algorithm **??**) with the following parameters:

i) $\epsilon_{\max}$: maximum allowed DoS variation within any cluster C_r (e.g., $\epsilon_{\max} = 2$ allows clusters of at most two levels of DoS separation),

ii) $\delta_{\max}$: maximum spatial diameter of a cluster, making sure that all arcs in the group are relatively close to each other,

iii) $\tau_{\mathcal{E}}^{(0)} := \frac{\left|\mathcal{E}^{\text{clust}} - \mathcal{E}^{\text{orig}}\right|}{\mathcal{E}^{\text{orig}}}$: initial tolerance on the relative exposure difference,

iv) $\tau_{\Phi}^{(0)} := \frac{\left|\Phi^{\text{clust}} - \Phi^{\text{orig}}\right|}{\Phi^{\text{orig}}}$: initial tolerance on the relative flow difference,

v) α_{relax}: relaxation factor for gradually enlarging $\tau_{\mathcal{E}}$ and τ_{Φ} if needed.

At each iteration, k-means is run in the spectral space with a modified objective that adds Lagrangian penalties λ_1, λ_2 for violating the exposure and flow constraints, while hard constraints enforce DoS homogeneity and bounded spatial diameter inside each cluster. If no feasible solution is found under the current tolerances, the algorithm increases $\tau_{\mathcal{E}}$ and τ_{Φ} by a factor $\alpha_{\text{relax}} > 1$ and retries. As a safety valve, if the constrained problem remains infeasible, the method falls back to standard unconstrained k-means in the same spectral embedding, ensuring that a clustering $\mathcal{C} = \{C_1, \ldots, C_k\}$ is always produced. In this way, the clustered network $\tilde{\mathcal{G}}^{\text{b}}$ approximates the original exposure and flow while strictly controlling intra-cluster safety heterogeneity and spatial extent, which is crucial for fair, safety-aware comparison across cities.

Algorithm 5 Constrained Spectral Clustering with Multi-Start and Outer Relaxation

Require: $\mathcal{G}^{\mathrm{b}} = (\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}})$, k, $\epsilon_{\max}$, $\delta_{\max}$, $\tau_{\mathcal{E}}$, τ_{Φ}, α_{relax}, $T_{\max}$, R, $J_{\max}$
Ensure: $\tilde{\mathcal{G}}^{\mathrm{b}} = (\mathcal{C}, \tilde{\mathcal{A}}^{\mathrm{b}}, \tilde{\mathcal{W}}^{\mathrm{b}})$

Build symmetric adjacency matrix A and compute the spectral embedding $x_u \in \mathbb{R}^m$ for all $u \in \mathcal{V}^{\mathrm{b}}$
best solution $\leftarrow \varnothing$
for $j = 1, \ldots, J_{\max}$ **do**
    best score $\leftarrow +\infty$
    for $r = 1, \ldots, R$ **do**
        Randomly initialize centroids $\{\mu_q\}_{q=1}^{k}$
        for $t = 1, \ldots, T_{\max}$ **do**
            Assign each $u \in \mathcal{V}^{\mathrm{b}}$ to the nearest cluster C_q satisfying

$$\max_{v \in C_q} \left| w^{\mathrm{b}}_{uv} - \min_{\ell \in C_q} w^{\mathrm{b}}_{u\ell} \right| \le \epsilon_{\max}, \qquad \max_{v \in C_q} d^{\mathrm{b}}_{uv} \le \delta_{\max}$$

            Update centroids $\mu_q \leftarrow \frac{1}{|C_q|} \sum_{u \in C_q} x_u$, for all $q = 1, \ldots, k$
            if cluster assignments do not change **then**
                break
            end if
        end for
        Build clustered graph $\tilde{\mathcal{G}}^{\mathrm{b}}$ and compute

$$\Delta\mathcal{E} = |\mathcal{E}^{\mathrm{orig}} - \mathcal{E}^{\mathrm{clust}}|, \qquad \Delta\Phi = |\Phi^{\mathrm{orig}} - \Phi^{\mathrm{clust}}|$$

        $S \leftarrow \Delta\mathcal{E}/\tau_{\mathcal{E}} + \Delta\Phi/\tau_{\Phi}$
        if $S <$ best score **then**
            Store current clustering as best solution
            best score $\leftarrow S$
            Store best $\Delta\mathcal{E}, \Delta\Phi$
        end if
    end for
    if best $\Delta\mathcal{E} \le \tau_{\mathcal{E}}$ and best $\Delta\Phi \le \tau_{\Phi}$ **then**
        return best solution
    end if
    $\tau_{\mathcal{E}} \leftarrow \alpha_{\mathrm{relax}} \cdot \tau_{\mathcal{E}}, \quad \tau_{\Phi} \leftarrow \alpha_{\mathrm{relax}} \cdot \tau_{\Phi}$
end for
return best solution

Since the clustering step is based on a k-means-type heuristic, the final solution may depend on the initial centroids. To reduce this sensitivity, the procedure is repeated over multiple random initializations and the best clustering is retained. In addition, the tolerance relaxation step is applied only after the inner clustering loop has converged. This separates the clustering phase from the feasibility-adjustment phase and avoids relaxing the tolerances before the algorithm has fully exploited the current admissible region. This modification is

expected to improve robustness and solution quality, at the cost of increased computational effort. In particular, multiple random initializations reduce sensitivity to poor local minima, while outer relaxation avoids premature acceptance of looser tolerances and therefore yields a more reliable assessment of whether the original tolerances are attainable.

7.3.1 Mathematical Guarantees

We start by demonstrating that the Laplacian defined, is indeed a Laplacian:

Theorem 7.1: (Legitimacy of the normalized Laplacian).

The matrix $\mathcal{L} = I - D^{-1/2}AD^{-1/2}$ is a valid symmetric normalized graph Laplacian.

Proof. The proof is shown in Appendix A. □

Parameter Determination

The parameters $\gamma > 0$ and $\alpha \geq 0$ in Eq. (7.22) control the influence of spatial distance and flow magnitude, respectively, on the adjacency weights. They must be chosen to ensure the resulting graph structure reflects meaningful connectivity for clustering. The proof of the Laplacian's properties holds for any $\gamma > 0$ and $\alpha \geq 0$. Therefore, the Laplacian $\mathcal{L}$ maintains its fundamental properties for any admissible choice of γ and α. These parameters should be calibrated based on the specific network and desired clustering outcome, for instance, by ensuring the eigenvalues of $\mathcal{L}$ exhibit a clear gap to suggest a natural number of clusters k.

Theorem 7.2: (Flow preservation under demand aggregation).

Let $\mathcal{G}^{\mathrm{b}} = (\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}})$ be a di-graph with OD demand $q_{(u,v)}$ and let $\mathcal{C} = \{C_1, \ldots, C_k\}$ be any partition of $\mathcal{V}^{\mathrm{b}}$ into non-empty clusters. Construct the clustered representation $\tilde{\mathcal{G}}^{\mathrm{b}} = (\mathcal{C}, \tilde{\mathcal{A}}^{\mathrm{b}}, \tilde{\mathcal{W}}^{\mathrm{b}})$ and define clustered OD demands and AoN clustered flows as in Section 7.2.1 (Eqs. (7.16)–(7.17)). Then the total network flow is preserved:

$$\Phi^{\mathrm{orig}} \triangleq \sum_{a \in \mathcal{A}^{\mathrm{b}}} \phi_a^{\mathrm{b}} = \sum_{\tilde{a} \in \tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\phi}_{\tilde{a}} = \Phi^{\mathrm{clust}}.$$

Proof. The proof is shown in Appendix A. □

7.4 Network Comparison

Our goal is to compare cycling-safety structure across cities in a way that remains meaningful when the underlying networks have different sizes. Let $\mathcal{H}^{\mathrm{b}} = \big(\mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}}), \mathcal{A}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}}), \mathcal{W}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})\big)$ be the smaller cycling network and let $\mathcal{G}^{\mathrm{b}}$ be a larger one. We cluster $\mathcal{G}^{\mathrm{b}}$ into the clustered-network $\tilde{\mathcal{G}}^{\mathrm{b}} = \big(\mathcal{C}, \tilde{\mathcal{A}}^{\mathrm{b}}, \tilde{\mathcal{W}}^{\mathrm{b}}\big)$ so that the number of clustered-nodes matches the number of nodes in $\mathcal{H}^{\mathrm{b}}$:

$$|\mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})| = |\mathcal{C}| =: n.$$

This size matching allows a fair comparison between $\tilde{\mathcal{G}}^{\mathrm{b}}$ and $\mathcal{H}^{\mathrm{b}}$. Throughout this section, we use the notation $\hat{\mu}^{\mathcal{H}^{\mathrm{b}}}$ and $\hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}$ to indicate which network a quantity belongs to.

The Gromov–Wasserstein distance treats exposure as a distribution over nodes and assesses whether both networks place their "exposure mass" in similar locations and with similar spatial relationships (e.g., comparable hotspots and overall spread), without requiring any explicit node matching. This helps identify cities with a similar global safety profile, such as risk concentrated in a few areas versus distributed across many corridors.

Formally, for each directed arc $(u, v) \in \mathcal{A}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})$, we define the cyclist exposure

$$\mathcal{E}_{uv}^{\mathcal{H}^{\mathrm{b}}} = \phi_{uv}^{\mathcal{H}^{\mathrm{b}}} \cdot f_{uv}^{\mathcal{H}^{\mathrm{b}}} \cdot d_{uv}^{\mathcal{H}^{\mathrm{b}}}, \tag{7.24}$$

where $\phi_{uv}^{\mathcal{H}^{\mathrm{b}}}$ is the cyclist flow on (u, v), $f_{uv}^{\mathcal{H}^{\mathrm{b}}}$ is the traffic factor, and $d_{uv}^{\mathcal{H}^{\mathrm{b}}}$ is the arc length. We assign an exposure "mass" to each node $u \in \mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})$ by summing the exposure over all incoming and outgoing arcs:

$$\mu^{\mathcal{H}^{\mathrm{b}}}(u) = \sum_{(u,v)\in\mathcal{A}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})} \mathcal{E}_{uv}^{\mathcal{H}^{\mathrm{b}}} + \sum_{(v,u)\in\mathcal{A}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})} \mathcal{E}_{vu}^{\mathcal{H}^{\mathrm{b}}}. \tag{7.25}$$

For the clustered network $\tilde{\mathcal{G}}^{\mathrm{b}}$, each directed clustered-arc $(C_r, C_s) \in \tilde{\mathcal{A}}^{\mathrm{b}}$ has an analogous inter-cluster exposure

$$\tilde{\mathcal{E}}_{rs} = \tilde{\phi}_{rs}^{\mathrm{inter}} \cdot \tilde{f}_{rs}^{\mathrm{inter}} \cdot \tilde{d}_{rs}^{\mathrm{inter}}, \tag{7.26}$$

where $\tilde{\phi}_{rs}^{\mathrm{inter}}$ is the aggregated inter-cluster flow, $\tilde{f}_{rs}^{\mathrm{inter}}$ is the inter-cluster traffic factor (derived from the inter-cluster DoS), and $\tilde{d}_{rs}^{\mathrm{inter}}$ is the flow-weighted inter-cluster distance. The mass of a clustered-node $C_r \in \mathcal{C}$ is defined as

$$\mu^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r) = \sum_{(C_r,C_s)\in\tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\mathcal{E}}_{rs} + \sum_{(C_s,C_r)\in\tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\mathcal{E}}_{sr} + \tilde{\mathcal{E}}_{C_r}^{\mathrm{intra}}, \tag{7.27}$$

where $\tilde{\mathcal{E}}_{C_r}^{\mathrm{intra}}$ denotes the intra-cluster exposure accumulated within cluster C_r.

We normalize these masses into probability measures:

$$\hat{\mu}^{\mathcal{H}^{\mathrm{b}}}(u) = \frac{\mu^{\mathcal{H}^{\mathrm{b}}}(u)}{\sum_{x\in\mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})} \mu^{\mathcal{H}^{\mathrm{b}}}(x)}, \qquad \hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r) = \frac{\mu^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r)}{\sum_{C_s\in\mathcal{C}} \mu^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_s)}. \tag{7.28}$$

These distributions describe how cyclist exposure is spatially allocated across each network.

To compare exposure structure without prescribing a node-to-node correspondence, we use the GW metric. This requires intra-network dissimilarity matrices. Let $x_u \in \mathbb{R}^2$ be the coordinates of node $u \in \mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})$ and let $c_{C_r} \in \mathbb{R}^2$ be the centroid of cluster $C_r \in \mathcal{C}$. We define

$$D^{\mathcal{H}^{\mathrm{b}}}(u,v) = \left|\hat{\mu}^{\mathcal{H}^{\mathrm{b}}}(u) - \hat{\mu}^{\mathcal{H}^{\mathrm{b}}}(v)\right| + \gamma \left\|x_u - x_v\right\|_2, \qquad \forall\, u,v \in \mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}}), \tag{7.29}$$

and

$$D^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r,C_s) = \left|\hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r) - \hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_s)\right| + \gamma \left\|c_{C_r} - c_{C_s}\right\|_2, \qquad \forall\, C_r, C_s \in \mathcal{C}, \tag{7.30}$$

where $\gamma > 0$ balances exposure dissimilarity and geographic separation.

Let $\Pi\big(\hat{\mu}^{\mathcal{H}^{\mathrm{b}}}, \hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}\big)$ be the set of admissible couplings $\pi \in \mathbb{R}_+^{n\times n}$ satisfying the marginal constraints

$$\sum_{C_s \in \mathcal{C}} \pi_{u,C_s} = \hat{\mu}^{\mathcal{H}^{\mathrm{b}}}(u), \qquad \sum_{u \in \mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}})} \pi_{u,C_s} = \hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_s). \tag{7.31}$$

The squared GW metric is then

$$GW^2(\mathcal{H}^{\mathrm{b}}, \tilde{\mathcal{G}}^{\mathrm{b}}) = \min_{\pi \in \Pi(\hat{\mu}^{\mathcal{H}^{\mathrm{b}}}, \hat{\mu}^{\tilde{\mathcal{G}}^{\mathrm{b}}})} \sum_{\substack{u,v \in \mathcal{V}^{\mathrm{b}}(\mathcal{H}^{\mathrm{b}}) \\ C_r, C_s \in \mathcal{C}}} \Big(D^{\mathcal{H}^{\mathrm{b}}}(u,v) - D^{\tilde{\mathcal{G}}^{\mathrm{b}}}(C_r,C_s)\Big)^2 \pi_{u,C_r}\, \pi_{v,C_s}, \tag{7.32}$$

and $GW(\mathcal{H}^{\mathrm{b}}, \tilde{\mathcal{G}}^{\mathrm{b}}) = \sqrt{GW^2(\mathcal{H}^{\mathrm{b}}, \tilde{\mathcal{G}}^{\mathrm{b}})}$. A small GW value indicates that the two networks have similar exposure organization (high-exposure regions and their spatial relationships) after clustering.

7.5 Benchmark Instances and Case Study

We illustrate the proposed comparison framework on the toy example in Fig. 7.2. The original small network is denoted by $\mathcal{H}^{\mathrm{b}}$, and the larger network $\mathcal{G}^{\mathrm{b}}$ is compressed into a clustered representation $\tilde{\mathcal{G}}^{\mathrm{b}}$ with $k = 4$ clusters, following the same clustering configuration as in Fig. 7.1. In this example, $\tilde{\mathcal{G}}^{\mathrm{b}}$ has roughly half as many nodes as $\mathcal{G}^{\mathrm{b}}$ while preserving the key safety attributes (flow, exposure, DoS) in aggregated form at the cluster level.

We compute the Wasserstein distance between the node–exposure measures of $\mathcal{H}^{\mathrm{b}}$ and $\tilde{\mathcal{G}}^{\mathrm{b}}$, obtaining $GW = 0.412$. In our normalized setting, this value indicates that the spatial allocation of cyclist exposure in the clustered network remains close to that of the original network, despite the substantial reduction in network size. The Wasserstein distance captures how total exposure mass is redistributed across space. In this benchmark, they jointly indicate that the clustered representation $\tilde{\mathcal{G}}^{\mathrm{b}}$ is a faithful and size-reduced surrogate of $\mathcal{H}^{\mathrm{b}}$ for safety comparison.

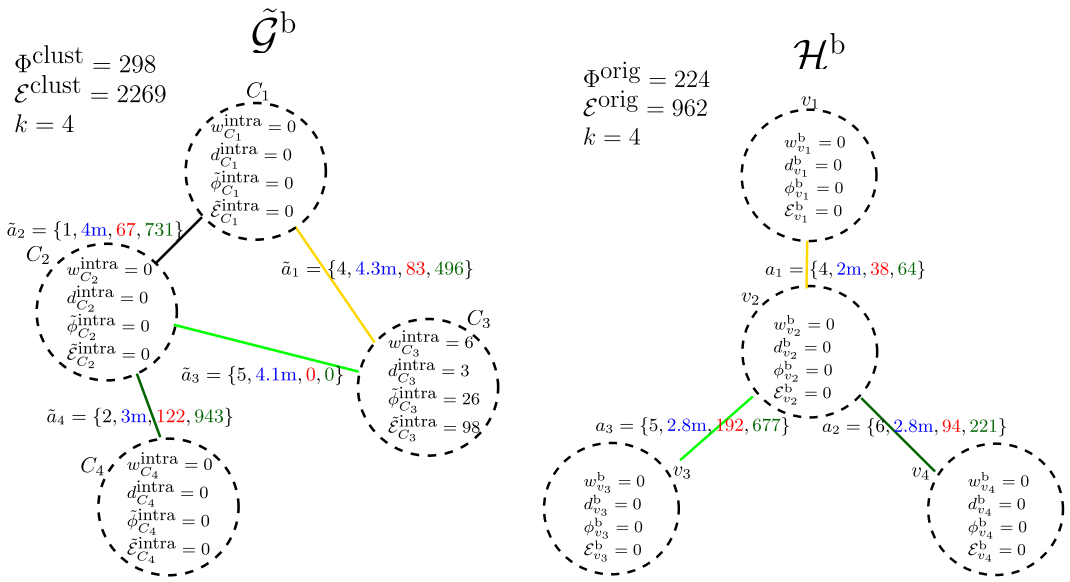

Figure 7.2: Academic example for clustering process

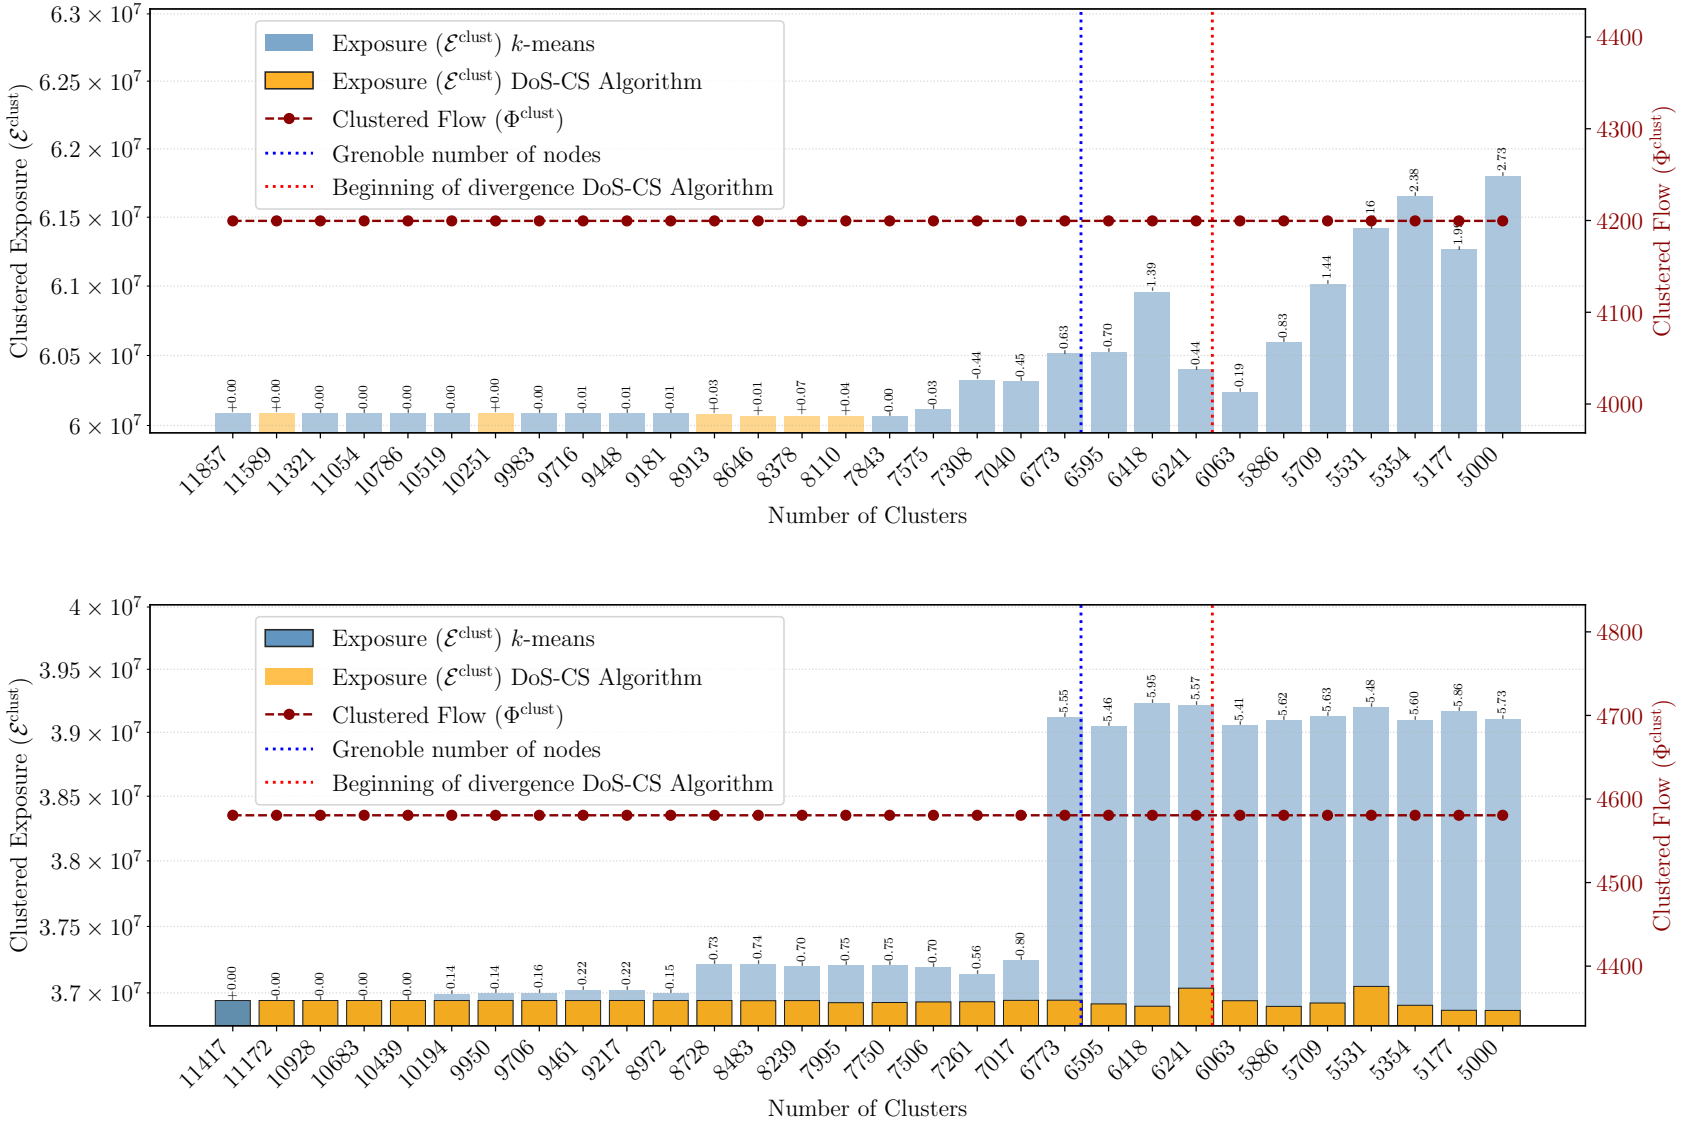

Figure 7.3: Evolution of cyclist exposure and cyclist flow through the DoS-CS algorithm (top: Lyon; bottom: Strasbourg).

Case Study: Grenoble, Lyon, Strasbourg

We illustrate the proposed framework on three French cities with well-developed cycling networks: Lyon (11 857 nodes), Strasbourg (11 417 nodes), and Grenoble (6 773 nodes). Grenoble is the smallest network and is taken as the reference. For each of the larger cities, we apply Algorithm **??** to cluster their network down to $|\mathcal{C}| = 6\,773$ clustered-nodes, yielding clustered networks $\tilde{\mathcal{G}}^{\mathrm{b}}_{\mathrm{Lyon}}$ and $\tilde{\mathcal{G}}^{\mathrm{b}}_{\mathrm{Strasbourg}}$ that are comparable in size to $\mathcal{H}^{\mathrm{b}}_{\mathrm{Grenoble}}$. We then compare each clustered network to Grenoble using the Wasserstein metricin Section 7.4.

Fig. 7.3 shows how the total clustered exposure $\mathcal{E}^{\mathrm{clust}}$ (bars, left axis) and total clustered flow Φ^{clust} (red dashed line, right axis) evolve as the number of clusters decreases, for Lyon (top) and Strasbourg (bottom). The blue vertical line marks the point where the number of clusters equals the number of Grenoble nodes (6 773), while the red vertical line indicates the onset of noticeable divergence between the original and clustered exposure. Two clustering strategies are compared: unconstrained k-means and Algorithm **??**. Across both cities, the DoS-CS algorithm consistently produces clustered networks whose total exposure remains much closer to the original value, particularly in the high-compression regime (towards the right of the plots). Lyon reaches the divergence threshold earlier: once the number of clusters falls below about 6 200, the exposure under k-means and (to a lesser extent) DoS-CS algorithm deviates more strongly from the original network. In contrast, Strasbourg maintains a near-constant exposure over a wider range of cluster counts, suggesting that its DoS structure can be compressed more aggressively without losing safety-relevant information. The clustered flow remains remarkably stable under Algorithm **??**, confirming that the total cyclist demand is well preserved during clustering.

Fig. 7.4 compares the spatial distribution of the degree of separation before and after clustering, for Lyon (top row) and Strasbourg (bottom row). The first column shows the original networks colored by DoS, the second the result of unconstrained k-means, and the third the result of Algorithm **??**. While both methods reduce the number of nodes to match Grenoble, the unconstrained k-means clustering tends to mix arcs of very different DoS within the same cluster, leading to patches where risky and safe segments are aggregated and the minimum DoS is not systematically preserved. Algorithm **??**, by construction, assigns to each cluster the worst (minimum) DoS of its constituent arcs and enforces bounds on intra-cluster DoS and spatial diameter. Visually, this yields clustered networks whose high-risk corridors and low-risk green-ways remain clearly aligned with those of the original networks, even under substantial node reduction. This qualitative behavior is consistent with the exposure and flow curves in Fig. 7.3.

The GW distance which compares the spatial organization of node-exposure masses through an optimal coupling yields a clear distinction:

$$GW(\mathcal{H}^{\mathrm{b}}_{\mathrm{Grenoble}}, \tilde{\mathcal{G}}^{\mathrm{b}}_{\mathrm{Strasbourg}}) = 0.0085,$$
$$GW(\mathcal{H}^{\mathrm{b}}_{\mathrm{Grenoble}}, \tilde{\mathcal{G}}^{\mathrm{b}}_{\mathrm{Lyon}}) = 0.0243.$$

Hence, GW indicates that Strasbourg's exposure distribution is substantially closer to Grenoble's than Lyon's is. Interestingly, this network-based similarity is coherent with independent

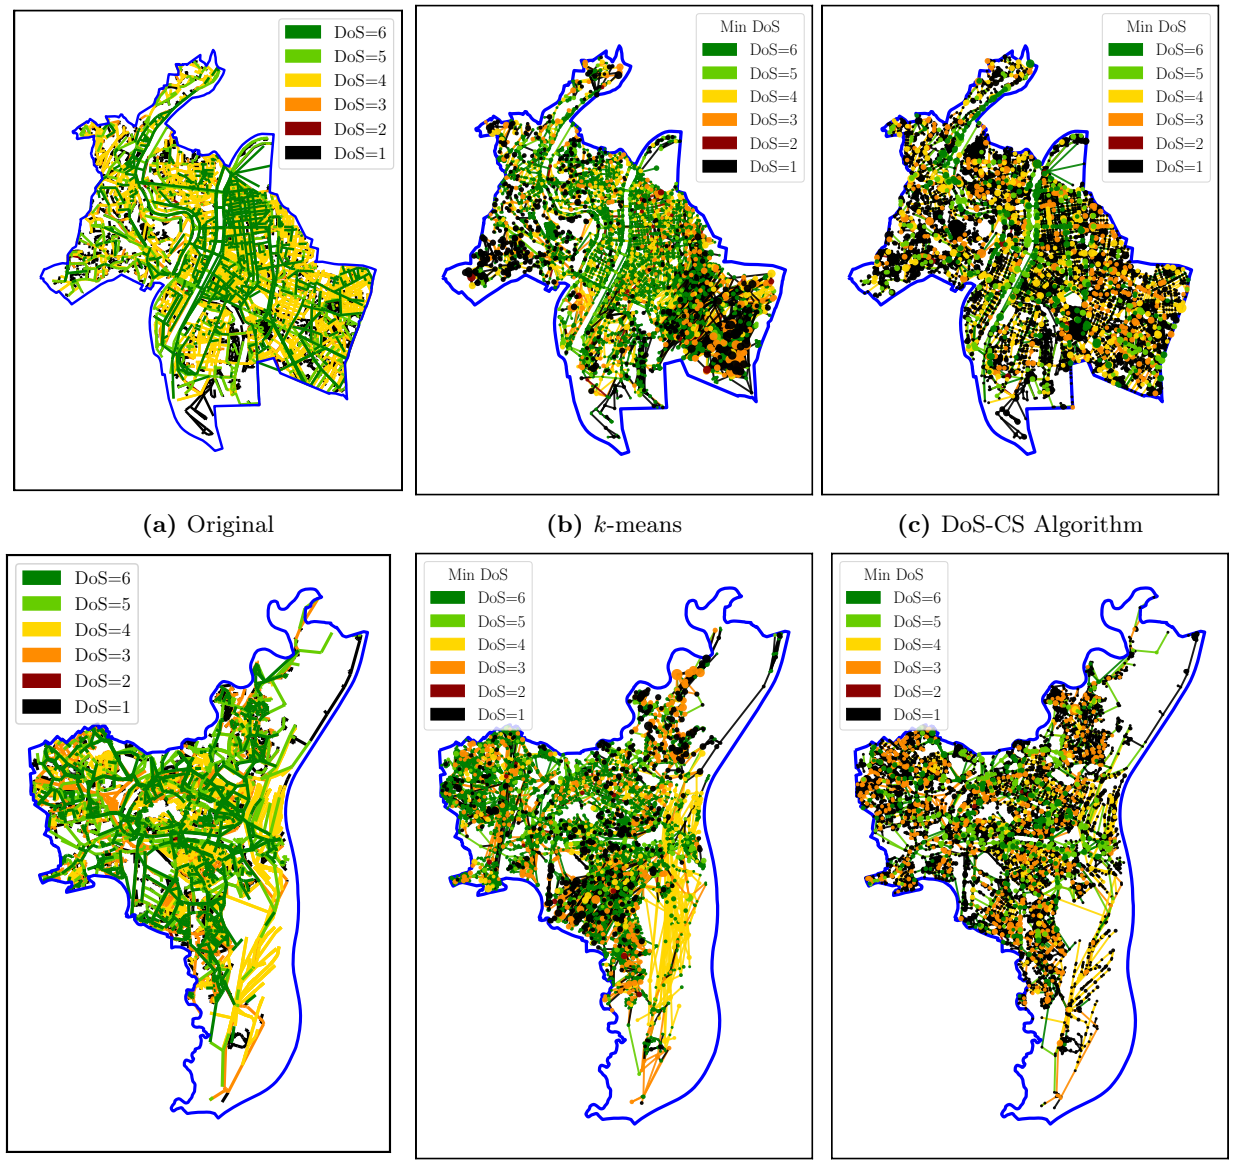

(a) Original **(b)** k-means **(c)** DoS-CS Algorithm

Figure 7.4: Clustering process with the DoS-CS algorithm. The first column shows the original networks, the second the k-means clustering, and the third the DoS-CS clustering (top: Lyon, bottom: Strasbourg).

survey data from the Baromètre Vélo project [103], where Grenoble and Strasbourg systematically rank among the most bicycle-friendly French cities in terms of perceived safety and comfort, while Lyon, although performing well, tends to score slightly lower. Our framework recovers this qualitative ordering using only infrastructure, flow, and exposure information, showing that clustering-based compression combined with safety-aware metrics can provide a fair and interpretable basis for cross-city comparisons of cycling safety.

7.6 Concluding Remarks

The present study focuses on benchmarking and comparison of existing networks. A natural next step is to use the clustered representation not only for analysis, but also for design and decision-making. In earlier work [53], we showed that the safety-oriented cycling network design problem is NP-hard, and we currently rely on heuristic methods (e.g., genetic algorithms) applied directly to the full graph. While effective, this strategy becomes computationally expensive as network size grows, which limits scalability and the ability to explore richer constraints (e.g., equity targets, multi-stage budgets, or multi-objective settings).

A promising direction is therefore to formulate a reduced version of the network design problem on the clustered graph $\tilde{\mathcal{G}}^{\mathrm{b}}$. At this level, the decision variables could represent cluster-to-cluster upgrades (e.g., improving the worst inter-cluster bottlenecks), intra-cluster reinforcement (e.g., ensuring safe connectivity within critical zones), or the prioritization of new protected corridors linking high-demand regions. An important methodological challenge will be to define principled lifting rules that map optimized cluster-level decisions back to implementable interventions on the original network. This can be evaluated by comparing the resulting exposure, flow, and fairness objectives with those obtained by optimizing directly on the full graph. More broadly, this opens the door to multi-resolution strategies in which optimization alternates between coarse (cluster) and fine (original) representations, allowing large instances to be handled at a fraction of the computational cost while still approaching high-quality solutions.

In addition to clustering-based reduction, an important methodological direction is to broaden the comparison layer. In this thesis we focus on similarity measures built from exposure and spatial structure, but different planning questions may call for different notions of similarity. Future work will therefore investigate alternative graph-comparison metrics that remain size-robust and policy-relevant, for example distances based on (i) safety-backbone connectivity (e.g., the structure of high-DoS corridors), (ii) accessibility or reachability to key destinations under low-stress constraints, (iii) distributional patterns across neighborhoods (equity-sensitive divergences), or (iv) multi-scale descriptors that combine local corridor structure with global network organization. These alternatives could complement the GW distance and help identify which metrics best correlate with external signals (e.g., surveys, crash proxies) and which are most useful for transferring design strategies across cities.

Beyond optimization speed-ups, several extensions can improve realism and policy relevance. First, integrating time-dependent demand and seasonal effects would allow the com-

parison metrics to capture robustness under temporal variability, while incorporating crash or near-miss data (when available) would strengthen the link between exposure-based indicators and observed safety outcomes. Second, the framework can be extended to multi-modal settings by coupling cycling networks with public transport and road networks, enabling comparisons that reflect accessibility rather than infrastructure alone. Third, alternative clustering objectives and OT-based distances could be tailored to specific planning questions, such as corridor prioritization, phased investment plans, or the identification of structurally critical "safety backbones" shared across cities.

Ultimately, the goal is to turn the proposed clustering and comparison methods into a scalable decision-support pipeline: a tool that allows cities to benchmark their cycling networks in a consistent way, identify comparable peer cities, and explore cost-effective, safety-improving interventions at metropolitan scale.

Chapter 8

Conclusions and Future Research

This thesis developed a sequence of decision-support models for improving cycling safety and comfort under budget constraints, with an emphasis on network-level effects. Across chapters, cycling infrastructure is represented as a weighted directed graph whose arc weights encode the Degree of Separation (DoS) where higher level indicate more segregated bike lanes from cars traffic offering more safety and comfort, and upgrades are modeled as discrete DoS increases. This common representation makes it possible to move progressively from topology-driven planning to flow-aware, multi-modal, and justice-aware optimization, and finally to fair cross-city benchmarking through safety-preserving network reduction.

Chapter 3 established the foundations of this approach by identifying graph-theory metrics that remain meaningful when safety is encoded as arc weights. A key methodological contribution was the monotonicity analysis: for upgrade-based planning to be well-posed, the chosen metrics should not degrade when infrastructure is made safer. This led to a set of interpretable and optimization-compatible metrics (including DoS-weighted node strength, an arc-based closeness centrality on the line graph, and path-connectivity measures), and to a practical link between network structure and perceived safety and comfort using survey data. On top of this predictive layer, a budget-constrained upgrade problem was formulated and solved with efficient heuristics, showing that connecting fragmented high-DoS paths can yield large perceived gains with limited investment, especially in networks where protected infrastructure exists but lacks continuity.

Chapter 4 then served as the transition from this purely structural perspective to a flow-aware network design framework. While topological metrics capture whether safe links form continuous routes, they do not take into account how many cyclists actually use each part of the network. This chapter illustrates, using a simple example, that topology alone can lead to misguided investment prioritization when structurally central links have little cycling demand. To bridge this gap, we retain the network-structure logic of the previous chapter but embed it in a demand-driven setting by introducing an all-or-nothing cyclist traffic assignment and a flow-weighted exposure objective. Minimizing exposure does not replace the perceived safety and comfort metrics introduced earlier; rather, it provides a complementary and operational proxy that remains consistent with the same underlying idea that safer networks are those that offer usable, connected low-stress routes, but now weighted by actual usage. In particular, the Grenoble case study shows that upgrade portfolios that reduce exposure also yield clear improvements in the largest connected component of high-DoS infrastructure, which was previously identified as a strong predictor of perceived safety and comfort. The planning problem therefore becomes a bi-level optimization problem sensitive to where cyclists travel and to how

upgrades shift route choices: upgrades change generalized route costs, route costs redistribute flows, and the redistributed flows determine exposure. The results on Grenoble highlighted that this flow-aware extension prioritizes different corridors than purely topological strategies, achieves larger reductions in aggregate exposure, and induces a meaningful redistribution of cyclists toward safer facilities.

Chapters 5 and 6 extended this flow-aware view from cycling-only evaluation to realistic system interactions. By coupling cyclist exposure with car travel time under capacity-sensitive user-equilibrium assignment, the multi-modal framework makes explicit the central trade-off faced by cities: safety improvements for cyclists can require reallocating road space, with potential congestion impacts for drivers. The results showed that substantial exposure reductions are often achievable with only modest car-travel-time increases when upgrades are targeted carefully. Building on this, the thesis incorporated induced modal shift and transport-justice objectives, moving beyond aggregate efficiency to distributional outcomes. The sufficientarian formulation in particular illustrated that equity-oriented planning can reduce the share of users below minimum safety and accessibility thresholds while maintaining competitive system-level performance, supporting the feasibility of justice-aware low-carbon planning.

Finally, Chapter 7 addressed a methodological limitation that arises in cross-city analysis: direct comparison of raw cycling graphs can be dominated by differences in network size and mapping resolution rather than genuine safety structure. The proposed constrained, safety-aware clustering compresses a large network into a fixed-size supernetwork while preserving safety-relevant attributes (DoS composition, distance, slope, traffic factors, flows, and exposure). On top of these reduced networks, an exposure-based similarity measure enable interpretable benchmarking across cities. Empirically, the reduced representations preserve exposure and flow much better than unconstrained baselines and retain enough structure to support consistent comparisons.

Overall, the thesis contributes (i) a monotonicity-grounded set of safety/comfort indicators suitable for upgrade optimization, (ii) flow-aware and exposure-based formulations that connect infrastructure decisions to demand and experienced risk, (iii) multi-modal and justice-aware bi-level optimization models that capture system feedback and distributional goals, and (iv) a safety-preserving reduction and comparison framework that enables fair benchmarking across cities. Taken together, these contributions reinforce a central message: cycling safety is not only a local attribute of individual links, but a network phenomenon shaped by connectivity, route options, demand patterns, multi-modal feedback, and equity considerations.

Future research directions

The work opens several directions to strengthen realism, scalability, and policy relevance. A key message from the discussion with the jury is that the next step should not only be additional behavioral or algorithmic complexity, but also a clearer accounting layer: translating the model's results into monetized and comparable metrics that can support multi-modal evaluation and quantify the relationship between efficiency and equity. In that spirit, several extensions are described below.

Toward monetized and comparable appraisal metrics (welfare / cost-benefit analysis baseline): While this thesis evaluates upgrades using exposure- and travel-time-based indicators (and their distribution across users and regions), a natural extension is to embed these outcomes in a welfare or cost-benefit framework. This would allow direct comparison across modes and interventions using a common unit (e.g., net present value), and would make it possible to report the "cost of equity" in efficiency terms. Concretely, changes in travel time, safety proxies, and cycling uptake could be monetized using standard appraisal ingredients (value of time, crash valuation, and health benefits such as HEAT-style valuation), and combined with investment and external costs. This direction aligns with welfare-optimal network-expansion approaches that explicitly account for induced demand and accumulated benefits over time, as in the spirit of Paulsen and Rich [61]. Importantly, such a baseline would complement (rather than replace) the engineering indicators used here by making the results more directly actionable for conventional transport decision-making.

Robustness, model complexity, and learning from comparisons: A complementary direction is to improve the robustness and transferability of the proposed framework. Although this thesis moves beyond ad hoc planning by using explicit objectives and constraints, several problems remain NP-hard and must be solved with heuristics in practice. A key warning, supported by recent work on robust bicycle-network design [160], is that adding more behavioral detail and more complex algorithms can make the results more sensitive to noise and uncertainty (e.g., imperfect demand estimates, OpenStreetMap errors, or long-term changes in travel patterns). This motivates work that models uncertainty explicitly and prioritizes upgrade plans that remain stable across perturbations and scenarios. In the same spirit, the cross-city comparison framework in Chapter 7 can be extended from benchmarking to learning for design. Instead of optimizing each city from scratch, a scalable alternative is to identify network patterns associated with "successful" cities (safer, easier to improve, more accessible) and use them to guide interventions elsewhere. This could be done by building datasets that link city outcomes (surveys, crashes, or evaluation indicators) to network features, and then training models that detect transferable design patterns. While collecting such labeled data is not trivial, this direction could support policy transfer and help prioritize investments when full optimization is too costly or impractical.

Improving and validating the DoS–CSCS calibration: In this thesis, DoS is the main and portable decision variable, and its effects already reflect many physical and traffic-related factors through the alignment with the CycleRAP Safety and Comfort Score (CSCS). Future work should therefore focus on validating and refining this DoS–CSCS calibration. Where data are available, we can test whether lower predicted exposure (and higher DoS) is actually associated with fewer crashes or higher perceived safety, using crash/injury records and/or survey evidence. This validation could also highlight situations where OSM tags are not sufficient, especially at intersections or short mixing zones suggesting targeted adjustments to the DoS classification or additional penalties at specific nodes. Beyond improving predictive credibility, such validation is also important for appraisal: it strengthens the link between the modeled indicators and monetizable safety outcomes.

Empirical validation of the exposure indicator: In Chapter 3, we assessed whether the proposed graph-theory metrics were consistent with perceived safety and comfort using the Baromètre Vélo data [103]. In contrast, Chapters 4—7 assume that cyclists' exposure to motor traffic is a meaningful proxy for unsafety and use it as the main optimization objective. While intuitive, this relationship has not been directly validated in our setting. Future work should therefore test the empirical relevance of the exposure measure by comparing exposure levels (at link, corridor, or neighborhood scale) with independent safety signals, such as perceived-danger reports from the Baromètre Vélo, recorded crash and injury data, and, where available, near-miss or conflict observations. Establishing (or refining) this relationship would strengthen the interpretability of the objective function and also improve the credibility of any subsequent welfare accounting based on exposure reductions.

From AoN cycling assignment to richer behavioral models: The all-or-nothing (AoN) assumption is computationally convenient and appropriate when congestion is negligible, but it can be refined to capture route diversity and preference heterogeneity. A natural extension is to represent cyclists as multiple user classes with different tolerances to traffic and infrastructure quality. For instance, Dill and McNeil [8] distinguish four broad groups: *strong and fearless*, *enthused and confident*, *interested but concerned*, and *no way, no how*. Incorporating such classes in the assignment (e.g., via stochastic user equilibrium or logit-based path choice with class-specific sensitivities to DoS, speed environment, and intersection penalties) would allow upgrades to be evaluated not only through shortest-path shifts, but also through changes in route diversity and, crucially, through the share of the population that becomes willing to cycle as the network improves. These richer behavioral layers can then feed the appraisal layer by producing more credible induced-demand and accessibility changes.

Dynamic and temporal robustness: Cycling and car demand vary by time of day, season, and weather, and congestion is inherently time-dependent. An important extension is therefore to move from static assignments to time-resolved models by incorporating time-sliced demand (e.g., morning peak, inter-peak, evening peak) and evaluating upgrade plans under multiple operating conditions. This can be approached using ideas from dynamic traffic assignment (DTA) [161] and extended to cycling in a similar spirit, following emerging work on dynamic or multi-period perspectives for bicycle planning [162, 163]. Embedding such temporal structure would enable robust formulations that select upgrades performing well across scenarios, and it would also provide the time profile needed for discounted benefit streams in an NPV-style evaluation.

Stronger multi-modal realism, additional modes, and transport-justice objectives: Future work can extend the multi-modal layer by explicitly coupling cycling with public transport (bike–transit access/egress), walking links, and constraints related to curb space (parking/loading) or charging/parking for emerging fleets. This would shift the evaluation from infrastructure quality alone to accessibility, enabling planners to quantify how upgrades change access to jobs, schools, and services across modes and user groups. A promising direction is

to embed these extensions within a transport-justice perspective [9, 10], while also reporting efficiency impacts relative to a welfare baseline so that distributional gains can be interpreted alongside their efficiency costs. On the car side, more detailed representations of capacity changes (e.g., signal timing, lane-use constraints, and curb-space effects of reallocating space to cycling) would improve realism and strengthen the interpretation of multi-modal trade-offs.

Scalability through multi-resolution optimization: Chapter 7 suggests a natural route to scalability: first solve a reduced design problem on the clustered graph $\tilde{\mathcal{G}}^{\mathrm{b}}$, using the same objectives considered in Chapters 4–7 (e.g., minimizing $\mathcal{E}^{\mathrm{clust}}$ and its variants), and then lift the resulting decisions back to the original network to evaluate the corresponding objective value $\mathcal{E}$. Beyond this one-shot lifting, a promising direction is a multi-resolution optimization loop that alternates between coarse (cluster-level) and fine (arc-level) decisions. Such a strategy could make it practical to handle larger cities, richer constraints, and multi-stage investment plans at a fraction of the computational cost. Importantly, this also supports appraisal at scale, because it enables repeated scenario evaluation (and thus robust and monetized comparisons) without prohibitive computation.

From optimization to planning practice: From a policy and planning perspective, a key limitation of this thesis is that it frames network upgrading as a well-defined optimization problem, while real cycling investments are shaped by governance, negotiation, and constraints that go beyond engineering performance. Future work should therefore strengthen the link between model outputs and decision-making by explicitly accounting for feasibility constraints such as street ownership, parking and freight needs, emergency access, construction disruption, and multi-year phasing. In parallel, evaluation should move beyond average performance and systematically document distributional impacts: which neighborhoods and user groups gain or lose in terms of exposure, accessibility, and car-delay, and whether proposed interventions reduce or reproduce existing inequalities. A welfare / cost-benefit analysis layer would complement this by translating impacts into comparable metrics for decision-makers, while still reporting distributional outcomes explicitly.

In summary, this thesis provides a coherent modeling pipeline for cycling-infrastructure planning: from monotone topological safety indicators to flow-aware, multi-modal, and justice-oriented optimization, and finally to fair cross-city benchmarking via safety-preserving reduction. A key next step is to connect these model outputs to monetized, comparable appraisal metrics to support cross-mode decisions and quantify the efficiency–equity trade-off. Building on this accounting layer, richer behavioral models, temporal robustness, and implementation constraints offer a path toward a comprehensive, policy-ready decision-support tool for safer and more equitable cycling networks.

Appendix A

Complementary proofs

Proofs and complements of Chapter 3

In Section 3.2.4 we introduced three global indicators $s_{\mathcal{G}}$, $c_{\mathcal{G}}$ and $l_{\mathcal{G}}$, built from node strength, arc-based closeness on the line graph, and the largest connected component of high-DoS corridors, respectively. Their key theoretical property is monotonicity with respect to upgrades of the DoS labels. For convenience, we restate this result here before providing the proofs.

Theorem 3.1 (Monotonicity of the proposed network metrics): Let $\mathcal{G}^{\mathrm{b}} = (\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}})$ be a fixed directed graph representing the cycling network and let $\boldsymbol{w} = [w_a^{\mathrm{b}}]_{a\in\mathcal{A}^{\mathrm{b}}}$ and $\tilde{\boldsymbol{w}} = [\tilde{w}_a^{\mathrm{b}}]_{a\in\mathcal{A}^{\mathrm{b}}}$ be two DoS assignments with

$$\tilde{w}_a^{\mathrm{b}} \geq w_a^{\mathrm{b}} \quad \text{for all } a \in \mathcal{A}^{\mathrm{b}}, \tag{A.1}$$

where $w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}} \in \{1, \ldots, 6\}$. Let $h : \{1, \ldots, 6\} \to \mathbb{R}_{>0}$ be a strictly increasing map from DoS levels to CSCS values, and let the effective arc costs used in shortest-path computations be of the form $\ell(a) = d_a^{\mathrm{b}} \cdot f(w_a^{\mathrm{b}})$ with $d_a^{\mathrm{b}} > 0$ and f strictly decreasing. Then the three global indicators satisfy

$$s_{\mathcal{G}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq s_{\mathcal{G}}(\boldsymbol{w}^{\mathrm{b}}), \qquad c_{\mathcal{G}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq c_{\mathcal{G}}(\boldsymbol{w}^{\mathrm{b}}), \qquad l_{\mathcal{G}}(\tilde{\boldsymbol{w}}^{\mathrm{b}}) \geq l_{\mathcal{G}}(\boldsymbol{w}^{\mathrm{b}}). \tag{A.2}$$

In particular, the normalized global strength, the arc-based closeness, and the largest connected-component metric are all monotone non-decreasing under any DoS upgrade $\boldsymbol{w} \mapsto \tilde{\boldsymbol{w}}$. We now prove each of these three inequalities in turn.

Proof. We use three elementary observations.

(A) Monotone sums. If $x_i \geq y_i$ for all i and all coefficients are nonnegative, then $\sum_i x_i \geq \sum_i y_i$. Moreover, dividing by a positive constant preserves inequalities.

(B) Shortest paths under cost decreases. If every arc cost in a graph weakly decreases, then every shortest-path distance weakly decreases: for any nodes u, v, $d_{\mathrm{sp}}^{\mathrm{new}}(u, v) \leq d_{\mathrm{sp}}^{\mathrm{old}}(u, v)$, since the cost of every u–v path decreases and the minimum over paths cannot increase.

(C) Nested feasible sets. If a set of admissible arcs expands, then the maximum of any nonnegative objective over connected subgraphs cannot decrease.

Proof of (i): global strength. For each node v, $s(v;\boldsymbol{w}) = \sum_{a\in\mathcal{N}_v} h(w_a^{\mathrm{b}})$. Because $\tilde{w}_a^{\mathrm{b}} \geq w_a^{\mathrm{b}}$ and h is increasing, we have $h(\tilde{w}_a^{\mathrm{b}}) \geq h(w_a^{\mathrm{b}})$ for all a. By (A), $s(v;\tilde{\boldsymbol{w}}) \geq s(v;\boldsymbol{w})$ for every v, hence the (averaged) global strength $s_{\mathcal{G}}(\tilde{\boldsymbol{w}}) \geq s_{\mathcal{G}}(\boldsymbol{w})$. Normalizing by the positive constant s_{FS} preserves the inequality, so $s_{\mathcal{G}}$ (and its normalized form) is monotone non-decreasing.

Proof of (ii): arc-based closeness on the line graph. In the line graph, each edge cost $\tilde{\ell}_{e_{a,a'}}(\boldsymbol{w})$ is built from positive lengths and a weight term that is non-increasing under DoS upgrades (because it uses $h(7-w)$ with h increasing). Therefore, upgrading DoS yields $\tilde{\ell}_e(\tilde{\boldsymbol{w}}) \leq \tilde{\ell}_e(\boldsymbol{w})$ for every line-graph edge e. By (B), all line-graph shortest-path distances weakly decrease: $\tilde{d}_{\tilde{\boldsymbol{w}}}(e_a, e_{a'}) \leq \tilde{d}_{\boldsymbol{w}}(e_a, e_{a'})$. Since closeness is inversely proportional to the sum of distances,

$$c(e_a;\boldsymbol{w}) = \frac{|\mathcal{A}^{\mathrm{b}}| - 1}{\sum_{e_{a'}} \tilde{d}_{\boldsymbol{w}}(e_a, e_{a'})},$$

the denominator weakly decreases, so $c(e_a;\tilde{\boldsymbol{w}}) \geq c(e_a;\boldsymbol{w})$ whenever the denominator is positive (and equality holds trivially in isolated cases by convention). Dividing by the positive constant $c_{FS}(e_a)$ and averaging over arcs preserves the inequality, hence $c_{\mathcal{G}}(\tilde{\boldsymbol{w}}) \geq c_{\mathcal{G}}(\boldsymbol{w})$.

Proof of (iii): largest connected-component corridor metric. Fix $\gamma \in \{2,\dots,6\}$ and define the high-DoS arc sets $C_\gamma(\boldsymbol{w}) = \{a : w_a^{\mathrm{b}} \geq \gamma\}$ and $C_\gamma(\tilde{\boldsymbol{w}}) = \{a : \tilde{w}_a^{\mathrm{b}} \geq \gamma\}$. Since $\tilde{w}_a^{\mathrm{b}} \geq w_a^{\mathrm{b}}$, we have $C_\gamma(\boldsymbol{w}) \subseteq C_\gamma(\tilde{\boldsymbol{w}})$. Thus the family of connected subgraphs available at threshold γ can only expand, so by (C) the maximized corridor score (sum of nonnegative terms $d_a^{\mathrm{b}} h(\cdot)$ over the corridor) cannot decrease. Normalizing by the fixed fully-safe constant preserves the inequality, hence $l_N(\gamma;\tilde{\boldsymbol{w}}) \geq l_N(\gamma;\boldsymbol{w})$ for all γ. Finally, $l_{\mathcal{G}}$ is a positive weighted average of these $l_N(\gamma;\cdot)$ values, so it is also monotone non-decreasing.

This proves $s_{\mathcal{G}}(\tilde{\boldsymbol{w}}) \geq s_{\mathcal{G}}(\boldsymbol{w})$, $c_{\mathcal{G}}(\tilde{\boldsymbol{w}}) \geq c_{\mathcal{G}}(\boldsymbol{w})$, and $l_{\mathcal{G}}(\tilde{\boldsymbol{w}}) \geq l_{\mathcal{G}}(\boldsymbol{w})$. □

Lemma 3.2 (Network-upgrade optimization problem complexity): Problem 3.2 in Eq. (3.38)–(3.40) is NP-hard. In particular, an exhaustive search over all admissible DoS configurations has worst-case complexity $\mathcal{O}(6^{|\mathcal{A}^{\mathrm{b}}|})$.

Proof. Recall that a feasible upgraded network configuration is described by a DoS assignment $\tilde{\boldsymbol{w}}^{\mathrm{b}} = [\tilde{w}_a^{\mathrm{b}}]_{a\in\mathcal{A}^{\mathrm{b}}} \in \Omega$, with

$$\Omega = \left\{ \tilde{\boldsymbol{w}}^{\mathrm{b}} \in \{1,\dots,6\}^{|\mathcal{A}^{\mathrm{b}}|} : w_a^{\mathrm{b}} \leq \tilde{w}_a^{\mathrm{b}} \leq 6,\ \forall a \in \mathcal{A}^{\mathrm{b}} \right\}, \tag{A.3}$$

and that the corresponding CSCS weights are given by $h(\tilde{w}_a^{\mathrm{b}})$, where $h : \{1,\dots,6\} \to \mathbb{R}_{>0}$ is strictly increasing. The objective of Problem 3.2,

$$\tilde{\boldsymbol{w}}^{\mathrm{b}^*} = \operatorname*{argmax}_{\tilde{\boldsymbol{w}}\in\Omega} \left\| (M_{\boldsymbol{w}^{\mathbf{b}}} - M_{\tilde{\boldsymbol{w}}^{\mathrm{b}}})\, \tilde{\theta} \right\| \tag{A.4}$$

represents the gain in predicted safety and comfort obtained by upgrading the DoS levels, subject to the budget constraint

$$\sum_{a\in\mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, \tilde{w}_a^{\mathrm{b}}) \ \leq\ B. \tag{A.5}$$

To connect this to a knapsack-type problem, introduce binary decision variables that encode the choice of upgraded DoS level for each arc. For each $a \in \mathcal{A}^{\mathrm{b}}$ and $k \in \{w_a^{\mathrm{b}}, \dots, 6\}$, define

$$x_{a,k} = \begin{cases} 1, & \text{if arc } a \text{ is assigned DoS level } k, \\ 0, & \text{otherwise.} \end{cases} \tag{A.6}$$

These variables satisfy

$$\sum_{k=w_a}^{6} x_{a,k} = 1 \quad \text{for all } a \in \mathcal{A}^{\mathrm{b}}, \tag{A.7}$$

and the upgraded DoS of arc a can be written as $\tilde{w}_a^{\mathrm{b}} = \sum_{k=w_a^{\mathrm{b}}}^{6} k \cdot x_{a,k}$, with CSCS weight $h(\tilde{w}_a^{\mathrm{b}}) = \sum_{k=w_a^{\mathrm{b}}}^{6} h(k) \cdot x_{a,k}$ by linearity.

The budget constraint can then be expressed as

$$\sum_{a \in \mathcal{A}^{\mathrm{b}}} \sum_{k=w_a^{\mathrm{b}}}^{6} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, k)\, x_{a,k} \;\le\; B, \tag{A.8}$$

which is a linear inequality in the binary variables $x_{a,k}$, where $d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, k)$ is the cost of upgrading arc a from its original DoS w_a^{b} to level k.

Because h is strictly increasing and the path-level metric $l_{\mathcal{G}_{\tilde{w}}^{\mathrm{b}}}$ is monotone non-decreasing with respect to the CSCS weights $h(\tilde{w}_a^{\mathrm{b}})$ (as shown in Section 3.2.4), the objective $\|(M_{\boldsymbol{w}^{\mathrm{b}}} - M_{\tilde{\boldsymbol{w}}^{\mathrm{b}}})\,\tilde{\theta}\|$ is a non-decreasing function of these upgrades. In particular, consider the restricted case where each arc can either remain at its original DoS level or be upgraded directly to the maximum level 6, i.e.

$$\tilde{w}_a^{\mathrm{b}} \in \{w_a^{\mathrm{b}}, 6\} \quad \Leftrightarrow \quad x_{a,w_a^{\mathrm{b}}}, x_{a,6} \in \{0,1\},\; x_{a,w_a^{\mathrm{b}}} + x_{a,6} = 1. \tag{A.9}$$

In this setting, each arc a corresponds to a single binary decision $x_{a,6} \in \{0,1\}$: upgrade a to DoS 6 ($x_{a,6} = 1$) or keep it at its original level ($x_{a,6} = 0$). The budget constraint reduces to

$$\sum_{a \in \mathcal{A}^{\mathrm{b}}} d_a^{\mathrm{b}} \cdot \pi(w_a^{\mathrm{b}}, 6) \cdot x_{a,6} \;\le\; B, \tag{A.10}$$

and the objective becomes a non-decreasing function of the set of upgraded arcs. By appropriate normalization of $l_{\mathcal{G}_{\tilde{w}}^{\mathrm{b}}}$ or of the linear model coefficients in $\tilde{\theta}$, the gain in the objective can be made proportional to a sum of per-arc "values" $v_a x_{a,6}$, where $v_a > 0$ represents the benefit of upgrading arc a.

Hence, in this restricted case, Problem 3.2 is equivalent to the classical 0–1 knapsack problem:

$$\max \sum_{a \in \mathcal{A}^{\mathrm{b}}} v_a \cdot x_{a,6} \quad \text{s.t.} \quad \sum_{a \in \mathcal{A}^{\mathrm{b}}} c_a \cdot x_{a,6} \le B, \quad x_{a,6} \in \{0,1\}, \tag{A.11}$$

with $c_a = d_a \cdot \pi(w_a, 6)$ and v_a the associated upgrade benefit. Since 0–1 knapsack is NP-hard, and Problem 3.2 contains it as a special case, Problem 3.2 is NP-hard as well.

Finally, note that each arc has at most 6 admissible DoS levels. A naive exhaustive search over all possible assignments $\tilde{w} \in \Omega$ would therefore require evaluating up to $6^{|\mathcal{A}^{\mathrm{b}}|}$ configurations in the worst case, which yields the stated worst-case complexity. □

Lemma 3.3 (MLCC Complexity): MLCC algorithm is $\mathcal{O}(|\mathcal{A}^{\mathrm{b}}||\mathcal{V}^{\mathrm{b}}|^4)$.

Proof. The more complex operations for MLCC algorithm involve computing the largest connected component and finding the shortest paths between them. Constructing a connected component utilizes the depth-first search algorithm, which has a computational complexity of $\mathcal{O}(|\mathcal{V}^{\mathrm{b}}| + |\mathcal{A}^{\mathrm{b}}|)$ [164]. However, for each component, the function traverses all nodes within the component, and for each node, it iterates over all its neighbors to verify whether they belong to the component according to their DoS. Each of these operations is executed at worst in $\mathcal{O}(|\mathcal{V}^{\mathrm{b}}|)$, then it would take at most $\mathcal{O}(|\mathcal{V}^{\mathrm{b}}|^3 + |\mathcal{V}^{\mathrm{b}}||\mathcal{A}^{\mathrm{b}}|)$. Finding the shortest paths is solved in $\mathcal{O}(|\mathcal{V}^{\mathrm{b}}|^2)$ because it uses Dijkstra's algorithm for weighted graphs [165], as in our case. Considering that in the worst scenario, each set of connected components contains $\frac{|\mathcal{V}^{\mathrm{b}}|}{2}$ elements and the process can occur at most $\frac{|\mathcal{A}^{\mathrm{b}}|}{2}$ times, the overall time complexity of this operation is at most $\mathcal{O}(|\mathcal{A}^{\mathrm{b}}||\mathcal{V}^{\mathrm{b}}|^4)$. Consequently, the overall time complexity of the function is dominated by the $\mathcal{O}(|\mathcal{A}^{\mathrm{b}}||\mathcal{V}^{\mathrm{b}}|^4)$ term. □

MLCC Flowchart

The flowchart in Fig. A.1 refers to the MLCC algorithm described in Chapter 3, keeping the same notation and meaning for $\hat{C}$, $C_{\tilde{\gamma}}$, and CC for connected component.

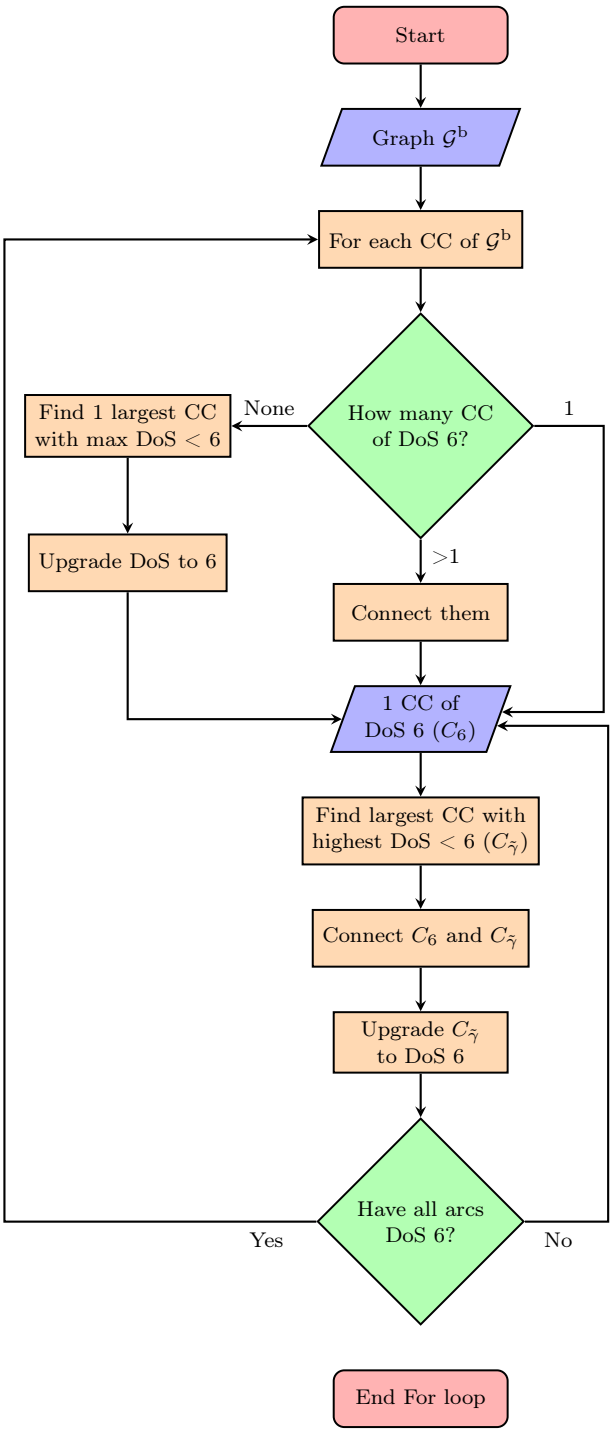

Figure A.1: Flowchart of the MLCC algorithm.

Proof of Chapter 5

Theorem 5.1 (Proof of Equivalence of UE): Consider the Beckmann potential associated with the car network,

$$\mathcal{B}^{\mathrm{c}}(\boldsymbol{\phi}^{\mathrm{c}}) = \sum_{a \in \mathcal{A}^{\mathrm{c}}} \int_0^{\phi_a^{\mathrm{c}}} t_a^{\mathrm{c}}(x)\,\mathrm{d}x,$$

where $t_a^{\mathrm{c}}(\cdot)$ is the travel-time function on arc a, and $\boldsymbol{\phi}^{\mathrm{c}}$ denotes the vector of arc flows.

Proof. Proof of i). For each arc $a \in \mathcal{A}^{\mathrm{c}}$, define

$$F_a(\phi_a^{\mathrm{c}}) := \int_0^{\phi_a^{\mathrm{c}}} t_a^{\mathrm{c}}(x)\,\mathrm{d}x.$$

Since the BPR travel-time function $t_a^{\mathrm{c}}(\cdot)$ is continuous and increasing, F_a is differentiable with

$$\frac{\mathrm{d}}{\mathrm{d}\phi_a^{\mathrm{c}}} F_a(\phi_a^{\mathrm{c}}) = t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}}).$$

Because $t_a^{\mathrm{c}}(\cdot)$ is increasing, the derivative of F_a is increasing, hence F_a is convex. If $t_a^{\mathrm{c}}(\cdot)$ is strictly increasing, then F_a is strictly convex. Therefore,

$$\mathcal{B}^{\mathrm{c}}(\boldsymbol{\phi}^{\mathrm{c}}) = \sum_{a \in \mathcal{A}^{\mathrm{c}}} F_a(\phi_a^{\mathrm{c}})$$

is convex as a sum of convex functions.

Proof of ii). Let X denote the feasible set defined by constraints Eq. (5.2)–(5.5). By construction, X is convex: flow-conservation constraints are linear, and nonnegativity constraints are convex. Suppose that $\boldsymbol{\phi}^{\mathrm{c}\star} \in X$ minimizes $\mathcal{B}^{\mathrm{c}}$ over X. Since $\mathcal{B}^{\mathrm{c}}$ is convex and differentiable, the first-order optimality condition for convex minimization gives

$$\nabla \mathcal{B}^{\mathrm{c}}(\boldsymbol{\phi}^{\mathrm{c}\star})^{\top} (\boldsymbol{\phi}^{\mathrm{c}} - \boldsymbol{\phi}^{\mathrm{c}\star}) \geq 0 \qquad \forall \boldsymbol{\phi}^{\mathrm{c}} \in X.$$

Now, for each arc a,

$$\frac{\partial \mathcal{B}^{\mathrm{c}}}{\partial \phi_a^{\mathrm{c}}}(\boldsymbol{\phi}^{\mathrm{c}\star}) = t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}\star}),$$

so the optimality condition becomes

$$\sum_{a \in \mathcal{A}^{\mathrm{c}}} t_a^{\mathrm{c}}(\phi_a^{\mathrm{c}\star})\,(\phi_a^{\mathrm{c}} - \phi_a^{\mathrm{c}\star}) \geq 0 \qquad \forall \boldsymbol{\phi}^{\mathrm{c}} \in X.$$

This is precisely the standard variational inequality formulation of the traffic assignment problem. It is well known to be equivalent to Wardrop's first principle: all used routes for a given OD pair have minimal travel time, and no traveler can unilaterally switch to a strictly shorter feasible route.

Therefore, every minimizer $\boldsymbol{\phi}^{\mathrm{c}\star}$ of the Beckmann potential is a user-equilibrium flow.

Finally, since the feasible set X is non-empty, closed, and convex, and since $\mathcal{B}^{\mathrm{c}}$ is continuous and convex, the minimization problem admits at least one solution.

□

Proofs of Chapter 7

Theorem 7.1 (Legitimacy of the normalized Laplacian): We consider the matrix defined in Chapter 7

$$\mathcal{L} \;=\; I - D^{-1/2} A D^{-1/2}, \tag{A.12}$$

where A is the (symmetric) weighted adjacency matrix and D is the diagonal degree matrix with $D_{uu} = \sum_v A_{uv}$. The goal of this proof is to show that $\mathcal{L}$ is the standard symmetric normalized graph Laplacian: it is symmetric by construction, and it satisfies the usual Laplacian properties used in spectral clustering (notably positive semidefiniteness and the classical spectral bounds). This justifies using the eigenvectors of $\mathcal{L}$ as a spectral embedding in our clustering framework.

Proof. The proof follows from the symmetry and non-negativity of the adjacency matrix A and the subsequent definition of $\mathcal{L}$.

i) Symmetry of A: Although $\mathcal{G}^{\mathrm{b}}$ is directed, we construct A on the induced *undirected* adjacency: for each unordered pair $\{u, v\}$ such that $(u, v) \in \mathcal{A}^{\mathrm{b}}$ or $(v, u) \in \mathcal{A}^{\mathrm{b}}$, we define symmetric arc attributes by

$$w^{\mathrm{b}}_{uv} := w^{\mathrm{b}}_{vu}, \qquad d^{\mathrm{b}}_{uv} := d^{\mathrm{b}}_{vu}, \qquad \phi^{\mathrm{sym}}_{uv} := \phi^{\mathrm{b}}_{(u,v)} + \phi^{\mathrm{b}}_{(v,u)}, \tag{A.13}$$

with the convention that $\phi^{\mathrm{b}}_{(u,v)} = 0$ if $(u, v) \notin \mathcal{A}^{\mathrm{b}}$ (and similarly for (v, u)). Since A_{uv} is defined as a nonnegative function of $(w^{\mathrm{b}}_{uv}, d^{\mathrm{b}}_{uv}, \phi^{\mathrm{sym}}_{uv})$ that is invariant under swapping u and v, it follows that $A_{uv} = A_{vu}$ for all u, v, and hence $A = A^{\top}$.

ii) Non-negativity of A: By construction, $w^{\mathrm{b}}_{uv} \geq 0$, $d^{\mathrm{b}}_{uv} \geq 0$, and $\phi^{\mathrm{sym}}_{uv} \geq 0$. Moreover, $\exp(-d^{\mathrm{b}}_{uv}/\gamma) > 0$ for $\gamma > 0$. Therefore, for every adjacent pair $\{u, v\}$, $A_{uv} \geq 0$. For non-adjacent nodes, $A_{uv} = 0$ by definition, so A is entrywise nonnegative.

iii) Properties of $\mathcal{L}$: Let $L = D - A$ be the (unnormalized) Laplacian. For any $\mathbf{y} \in \mathbb{R}^{|\mathcal{V}^{\mathrm{b}}|}$,

$$\mathbf{y}^{\top} L \mathbf{y} = \mathbf{y}^{\top}(D - A)\mathbf{y} = \frac{1}{2} \sum_{u,v} A_{uv} \, (y_u - y_v)^2 \;\geq\; 0, \tag{A.14}$$

so L is positive semidefinite. Since $A = A^{\top}$, we also have $D = D^{\top}$ and hence $L = L^{\top}$.

Assume $D_{uu} > 0$ for all u (isolated nodes can be removed or handled separately in the clustering step). Then the normalized Laplacian satisfies

$$\mathcal{L} = I - D^{-1/2} A D^{-1/2} = D^{-1/2}(D - A)D^{-1/2} = D^{-1/2} L D^{-1/2}. \tag{A.15}$$

For any $\mathbf{x}$, let $\mathbf{y} = D^{-1/2}\mathbf{x}$. Then

$$\mathbf{x}^\top \mathcal{L}\,\mathbf{x} = (D^{-1/2}\mathbf{x})^\top L(D^{-1/2}\mathbf{x}) = \mathbf{y}^\top L\mathbf{y} \;\geq\; 0, \tag{A.16}$$

which shows that $\mathcal{L}$ is positive semidefinite. Finally, since A and $D^{-1/2}$ are symmetric, $\mathcal{L}$ is symmetric as well, and therefore all eigenvalues of $\mathcal{L}$ are real and nonnegative. Moreover, by standard results for the symmetric normalized Laplacian, its spectrum is contained in $[0, 2]$. This completes the proof.

□

Theorem 7.1 (Flow Preservation in Network Clustering): Let $\mathcal{G}^{\mathrm{b}} = (\mathcal{V}^{\mathrm{b}}, \mathcal{A}^{\mathrm{b}}, \mathcal{W}^{\mathrm{b}})$ be the original cycling digraph with OD demand matrix $(q^{\mathrm{b}}_{(u,v)})$, and let $\tilde{\mathcal{G}}^{\mathrm{b}} = (\tilde{\mathcal{V}}^{\mathrm{b}}, \tilde{\mathcal{A}}^{\mathrm{b}}, \tilde{\mathcal{W}}^{\mathrm{b}})$ be its clustered (super-network) representation produced by Algorithm **??**. The key point is that clustering only aggregates origins/destinations and routes at the cluster level: it does not create or remove demand. As a result, the all-or-nothing assignment on $\tilde{\mathcal{G}}^{\mathrm{b}}$ redistributes exactly the same total OD mass as in the original network, but through super-arcs instead of original arcs. Therefore the total assigned flow is preserved, i.e.,

$$\Phi^{\mathrm{orig}} \triangleq \sum_{a\in\mathcal{A}^{\mathrm{b}}} \phi^{\mathrm{b}}_a \;=\; \sum_{\tilde{a}\in\tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\phi}_{\tilde{a}} \;=\; \Phi^{\mathrm{clust}}. \tag{A.17}$$

Proof. Observe that the clustered demand $\tilde{q}_{(C_r,C_{r'})}$ in Eq. (7.16) aggregates all original demands between nodes in clusters C_r and $C_{r'}$ as

$$\tilde{q}_{(C_r,C_{r'})} = \sum_{\substack{u\in C_r\\ v\in C_{r'}}} q^{\mathrm{b}}_{(u,v)}, \tag{A.18}$$

ensuring total demand preservation

$$\sum_{u,v\in\mathcal{V}^{\mathrm{b}}} q^{\mathrm{b}}_{(u,v)} = \sum_{C_r,C_{r'}\in\tilde{\mathcal{V}}^{\mathrm{b}}} \tilde{q}_{(C_r,C_{r'})}. \tag{A.19}$$

In the original graph $\mathcal{G}^{\mathrm{b}}$, the all-or-nothing assignment Eq. (4.4) yields arc flows

$$\phi^{\mathrm{b}}_a = \sum_{u,v} q^{\mathrm{b}}_{(u,v)} \cdot \mathbb{I}_{a\in p^{\mathrm{b}}_{(u,v)}}, \tag{A.20}$$

while in the clustered graph $\tilde{\mathcal{G}}^{\mathrm{b}}$, Eq. (7.17) gives superarc flows

$$\tilde{\phi}_{\tilde{a}} = \sum_{C_r,C_{r'}} \tilde{q}_{(C_r,C_{r'})} \cdot \mathbb{I}_{\tilde{a}\in\tilde{p}_{(C_r,C_{r'})}}. \tag{A.21}$$

The total flow in $\mathcal{G}^{\mathrm{b}}$ decomposes as:

$$\Phi^{\mathrm{orig}} = \sum_{a\in\mathcal{A}^{\mathrm{b}}} \phi^{\mathrm{b}}_a = \sum_{u,v} q^{\mathrm{b}}_{(u,v)} \underbrace{\sum_{a\in\mathcal{A}^{\mathrm{b}}} \mathbb{I}_{a\in p^{\mathrm{b}}_{(u,v)}}}_{=1} \tag{A.22}$$

$$= \sum_{u,v} q^{\mathrm{b}}_{(u,v)}, \tag{A.23}$$

since each path $p_{(u,v)}$ covers exactly one set of arcs. Similarly, for $\tilde{\mathcal{G}^{\mathrm{b}}}$:

$$\Phi^{\mathrm{clust}} = \sum_{\tilde{a}\in\tilde{\mathcal{A}}^{\mathrm{b}}} \tilde{\phi}_{\tilde{a}} = \sum_{C_r,C_{r'}} \tilde{q}_{(C_r,C_{r'})} \underbrace{\sum_{\tilde{a}\in\tilde{\mathcal{A}}^{\mathrm{b}}} \mathbb{I}_{\tilde{a}\in\tilde{p}_{(C_r,C_{r'})}}}_{=1} \tag{A.24}$$

$$= \sum_{C_r,C_{r'}} \tilde{q}_{(C_r,C_{r'})}, \tag{A.25}$$

where the indicator sums equal 1 because each clustered path $\tilde{p}_{(C_r,C_{r'})}$ traverses exactly one sequence of superarcs. By demand conservation, we conclude $\Phi^{\mathrm{orig}} \equiv \Phi^{\mathrm{clust}}$. □

Bibliography

[1] Jan Garrard, Chris Rissel, and Adrian Bauman. "Health benefits of cycling." In: *City cycling* 31 (2012), pp. 31–56 (cit. on pp. iii, 35).

[2] Elliot Fishman, Simon Washington, and Narelle Haworth. "Barriers and facilitators to public bicycle scheme use: A qualitative approach." In: *Transportation research part F: traffic psychology and behaviour* 15.6 (2012), pp. 686–698 (cit. on pp. iii, 5, 35, 71, 85).

[3] Joseph Broach, Jennifer Dill, and John Gliebe. "Where do cyclists ride? A route choice model developed with revealed preference GPS data." In: *Transportation Research Part A: Policy and Practice* 46.10 (2012), pp. 1730–1740 (cit. on pp. iii, 6, 35, 71).

[4] Brendan M Lawrence et al. "Cyclist exposure to the risk of car door collisions in mixed function activity centers: A study in Melbourne, Australia." In: *Traffic injury prevention* 19.sup1 (2018), S164–S168 (cit. on pp. iii, 6, 35, 71, 85, 104).

[5] International Road Assessment Programme. *CycleRAP*. Accessed: 2024-07-10. 2024 (cit. on pp. iii, 10, 30, 36, 74).

[6] Francisco Aparecido Rodrigues. "Network centrality: an introduction." In: *A mathematical modeling approach from nonlinear dynamics to complex systems*. Springer, 2018, pp. 177–196 (cit. on pp. iv, 36, 39, 41).

[7] Karel Martens. *Transport justice: Designing fair transportation systems*. Routledge, 2017 (cit. on pp. v, 18, 104, 106, 115).

[8] Jennifer Dill and Nathan McNeil. "Four types of cyclists? Examination of typology for better understanding of bicycling behavior and potential." In: *Transportation Research Record* 2387.1 (2013), pp. 129–138 (cit. on pp. vii, 7, 154).

[9] Mauro Salazar et al. "On Accessibility Fairness in Intermodal Autonomous Mobility-on-Demand Systems." In: *IFAC-PapersOnLine* 58.10 (2024), pp. 327–333 (cit. on pp. viii, 18, 106, 115, 155).

[10] Mauro Salazar et al. "Mobilizing transport justice: a sufficientarian optimization framework for intermodal mobility systems." In: *npj Sustainable Mobility and Transport* 2.1 (2025), p. 43 (cit. on pp. viii, 106, 155).

[11] Pekka Oja et al. "Health benefits of cycling: a systematic review." In: *Scandinavian journal of medicine & science in sports* 21.4 (2011), pp. 496–509 (cit. on p. 2).

[12] Thomas Götschi, Jan Garrard, and Billie Giles-Corti. "Cycling as a part of daily life: a review of health perspectives." In: *Transport Reviews* 36.1 (2016), pp. 45–71 (cit. on p. 2).

[13] A Wagner et al. "Leisure-time physical activity and regular walking or cycling to work are associated with adiposity and 5 y weight gain in middle-aged men: the PRIME Study." In: *International journal of obesity* 25.7 (2001), pp. 940–948 (cit. on p. 2).

[14] Gang Hu et al. "Occupational, commuting, and leisure-time physical activity in relation to risk for Type 2 diabetes in middle-aged Finnish men and women." In: *Diabetologia* 46.3 (2003), pp. 322–329 (cit. on p. 2).

[15] Noriyuki Nakanishi and Kenji Suzuki. "Daily life activity and the risk of developing hypertension in middle-aged Japanese men." In: *Archives of internal medicine* 165.2 (2005), pp. 214–220 (cit. on p. 2).

[16] Noël C Barengo et al. "Different forms of physical activity and cardiovascular risk factors among 24–64-year-old men and women in Finland." In: *European Journal of Cardiovascular Prevention & Rehabilitation* 13.1 (2006), pp. 51–59 (cit. on p. 2).

[17] Patrik Wennberg et al. "The effects of commuting activity and occupational and leisure time physical activity on risk of myocardial infarction." In: *European Journal of Preventive Cardiology* 13.6 (2006), pp. 924–930 (cit. on p. 2).

[18] Karim Abu-Omar and Alfred Rütten. "Relation of leisure time, occupational, domestic, and commuting physical activity to health indicators in Europe." In: *Preventive medicine* 47.3 (2008), pp. 319–323 (cit. on p. 2).

[19] David R Bassett et al. "Walking, cycling, and obesity rates in Europe, North America, and Australia." In: *Journal of physical activity and health* 5.6 (2008), pp. 795–814 (cit. on p. 2).

[20] Martin Lindström. "Means of transportation to work and overweight and obesity: a population-based study in southern Sweden." In: *Preventive medicine* 46.1 (2008), pp. 22–28 (cit. on p. 2).

[21] Penny Gordon-Larsen et al. "Active commuting and cardiovascular disease risk: the CARDIA study." In: *Archives of internal medicine* 169.13 (2009), pp. 1216–1223 (cit. on p. 2).

[22] Natalie Mueller et al. "Health impact assessment of active transportation: a systematic review." In: *Preventive medicine* 76 (2015), pp. 103–114 (cit. on p. 2).

[23] David Rojas-Rueda et al. "Health impact assessment of increasing public transport and cycling use in Barcelona: a morbidity and burden of disease approach." In: *Preventive medicine* 57.5 (2013), pp. 573–579 (cit. on p. 2).

[24] James Woodcock, Moshe Givoni, and Andrei Scott Morgan. "Health impact modelling of active travel visions for England and Wales using an Integrated Transport and Health Impact Modelling Tool (ITHIM)." In: *PloS one* 8.1 (2013), e51462 (cit. on p. 3).

[25] Todd M Litman. "Evaluating transportation equity: Guidance for incorporating distributional impacts in transport planning." In: *Institute of Transportation Engineers. ITE Journal* 92.4 (2022), pp. 43–49 (cit. on p. 3).

[26] Margaret Kent and Alex Karner. "Prioritizing low-stress and equitable bicycle networks using neighborhood-based accessibility measures." In: *International journal of sustainable transportation* 13.2 (2019), pp. 100–110 (cit. on p. 3).

[27] Luyu Liu, Jinhyung Lee, and Harvey J Miller. "Evaluating accessibility benefits and ridership of bike-transit integration through a social equity lens." In: *Computers, Environment and Urban Systems* 112 (2024), p. 102150 (cit. on pp. 3, 4).

[28] Zheyan Chen, Dea Van Lierop, and Dick Ettema. "Perceived accessibility: How access to dockless bike-sharing impacts activity participation." In: *Travel Behaviour and Society* 27 (2022), pp. 128–138 (cit. on p. 4).

[29] Sonja Kahlmeier et al. "Health economic assessment tools (HEAT) for walking and for cycling." In: (2011) (cit. on p. 4).

[30] Gerard Deenihan and Brian Caulfield. "Estimating the health economic benefits of cycling." In: *Journal of Transport & Health* 1.2 (2014), pp. 141–149 (cit. on p. 4).

[31] Elliot Fishman, Paul Schepers, and Carlijn Barbara Maria Kamphuis. "Dutch cycling: quantifying the health and related economic benefits." In: *American journal of public health* 105.8 (2015), e13–e15 (cit. on p. 4).

[32] Stefan Gössling and Andy S Choi. "Transport transitions in Copenhagen: Comparing the cost of cars and bicycles." In: *Ecological economics* 113 (2015), pp. 106–113 (cit. on p. 4).

[33] Fiona Crawford and Robin Lovelace. "The benefits of getting England cycling." In: (2015) (cit. on pp. 4, 5).

[34] Thomas Blondiau, Bruno Van Zeebroeck, and Holger Haubold. "Economic benefits of increased cycling." In: *Transportation research procedia* 14 (2016), pp. 2306–2313 (cit. on p. 5).

[35] Oliver Tristan Mytton, Jenna Panter, and David Ogilvie. "Longitudinal associations of active commuting with wellbeing and sickness absence." In: *Preventive medicine* 84 (2016), pp. 19–26 (cit. on p. 5).

[36] Ingrid JM Hendriksen et al. "The association between commuter cycling and sickness absence." In: *Preventive medicine* 51.2 (2010), pp. 132–135 (cit. on p. 5).

[37] Jennifer Dill and Kim Voros. "Factors affecting bicycling demand: initial survey findings from the Portland, Oregon, region." In: *Transportation research record* 2031.1 (2007), pp. 9–17 (cit. on p. 5).

[38] Jacob Larsen, Ahmed El-Geneidy, and Farhana Yasmin. "Beyond the quarter mile: Re-examining travel distances by active transportation." In: *Canadian Journal of Urban Research* 19.1 (2010), pp. 70–88 (cit. on p. 6).

[39] Andrew T Kaczynski, Melissa J Bopp, and Pamela Wittman. "Association of workplace supports with active commuting." In: *Preventing chronic disease* 7.6 (2010), A127 (cit. on p. 6).

[40] Ryan Swiers, Catherine Pritchard, and Ivan Gee. "A cross sectional survey of attitudes, behaviours, barriers and motivators to cycling in University students." In: *Journal of Transport & Health* 6 (2017), pp. 379–385 (cit. on p. 6).

[41] Javier Molina-García, Isabel Castillo, and James F Sallis. "Psychosocial and environmental correlates of active commuting for university students." In: *Preventive medicine* 51.2 (2010), pp. 136–138 (cit. on p. 6).

[42] Adriana A De Souza, Suely P Sanches, and Marcos AG Ferreira. "Influence of attitudes with respect to cycling on the perception of existing barriers for using this mode of transport for commuting." In: *Procedia-Social and Behavioral Sciences* 162 (2014), pp. 111–120 (cit. on p. 6).

[43] Katarzyna Iwińska et al. "Cycling in Warsaw, Poland–Perceived enablers and barriers according to cyclists and non-cyclists." In: *Transportation research part A: policy and practice* 113 (2018), pp. 291–301 (cit. on p. 6).

[44] John Pucher et al. "Walking and cycling in the United States, 2001–2009: evidence from the National Household Travel Surveys." In: *American journal of public health* 101.S1 (2011), S310–S317 (cit. on pp. 6, 85, 104).

[45] John Pucher and Ralph Buehler. "Cycling towards a more sustainable transport future." In: *Transport reviews* 37.6 (2017), pp. 689–694 (cit. on p. 7).

[46] J Davis. "Bicycle Safety Evaluation. Auburn University, City of Chattanooga, and Chattanooga-Hamilton County Regional Planning Commission, Chattanooga, Tenn., 1987. DELOITTE. Guia Para Melhorar a Governança Corporativa Através de Eficazes Controles Internos. Deloit." In: *Outubro* (1987), pp. 1002–1438 (cit. on pp. 7, 28).

[47] Alex Sorton and Thomas Walsh. "Bicycle stress level as a tool to evaluate urban and suburban bicycle compatibility." In: *Transportation Research Record* (1994), pp. 17–17 (cit. on pp. 8, 28).

[48] Bruce W Landis. *Bicycle interaction hazard score: a theoretical model*. Tech. rep. The National Academies of Sciences, Engineering, and Medicine, 1994 (cit. on pp. 8, 28).

[49] Shawn Turner, C Scott Shafer, William P Stewart, et al. *Bicycle suitability criteria for state roadways in Texas*. Tech. rep. Texas Transportation Institute, 1997 (cit. on pp. 8, 28).

[50] David L Harkey, Donald W Reinfurt, Alex Sorton, et al. *The Bicycle Compatibility Index: A level of service concept, implementation manual*. Tech. rep. United States. Federal Highway Administration, 1998 (cit. on pp. 8, 28).

[51] Highway Capacity Manual. "Highway capacity manual." In: *Washington, DC* 2.1 (2000), p. 1 (cit. on pp. 9, 28).

[52] Peter G Furth, Maaza C Mekuria, and Hilary Nixon. "Network connectivity for low-stress bicycling." In: *Transportation research record* 2587.1 (2016), pp. 41–49 (cit. on p. 9).

[53] Manuel Campero Jurado, Carlos Canudas-de Wit, and Giovanni De Nunzio. "A Graph-Based Approach to the Topological Optimization of Cycling Networks for the Improvement of Safety and Comfort of Cyclists." In: *IEEE Transactions on Intelligent Transportation Systems* (2025) (cit. on pp. 9, 93, 114, 149).

[54] Manuel Campero Jurado et al. "A General Graph-based Optimization Framework for Improving Safety and Comfort in Cycling Networks." In: *Proceedings of the 27th IEEE International Conference on Intelligent Transportation Systems*. Accepted for publication. 2024 (cit. on pp. 9, 14, 16, 37, 105).

[55] Jonathan Nolan, James Sinclair, and Jim Savage. "Are bicycle lanes effective? The relationship between passing distance and road characteristics." In: *Accident Analysis & Prevention* 159 (2021), p. 106184 (cit. on p. 9).

[56] Mahmoud Mesbah, Russell Thompson, and Sara Moridpour. "Bilevel optimization approach to design of network of bike lanes." In: *Transportation research record* 2284.1 (2012), pp. 21–28 (cit. on pp. 13, 16, 30, 37, 73, 86, 105).

[57] Heather Lauren Smith and Ali Haghani. *A mathematical optimization model for a bicycle network design considering bicycle level of service*. Tech. rep. University of Maryland, College Park, 2012 (cit. on pp. 13, 37, 73, 86, 105).

[58] Jen-Jia Lin and Chia-Jung Yu. "A bikeway network design model for urban areas." In: *Transportation* 40 (2013), pp. 45–68 (cit. on pp. 13, 16).

[59] Antonio Mauttone et al. "Bicycle network design: model and solution algorithm." In: *Transportation research procedia* 27 (2017), pp. 969–976 (cit. on pp. 13, 16, 37, 73).

[60] Jennifer Duthie and Avinash Unnikrishnan. "Optimization framework for bicycle network design." In: *Journal of Transportation engineering* 140.7 (2014), p. 04014028 (cit. on pp. 13, 16, 37, 73, 87, 105).

[61] Mads Paulsen and Jeppe Rich. "Societally optimal expansion of bicycle networks." In: *Transportation Research Part B: Methodological* 174 (2023), p. 102778 (cit. on pp. 13, 16, 37, 73, 87, 105, 129, 153).

[62] Nina Wiedemann et al. "Bike network planning in limited urban space." In: *Transportation Research Part B: Methodological* 192 (2025), p. 103135 (cit. on pp. 13, 16, 86, 91, 105).

[63] Luis Guillermo Natera Orozco et al. "Data-driven strategies for optimal bicycle network growth." In: *Royal society open science* 7.12 (2020), p. 201130 (cit. on pp. 13, 16, 133).

[64] Luis E Olmos et al. "A data science framework for planning the growth of bicycle infrastructures." In: *Transportation research part C: emerging technologies* 115 (2020), p. 102640 (cit. on pp. 14, 16).

[65] Christoph Steinacker et al. "Demand-driven design of bicycle infrastructure networks for improved urban bikeability." In: *Nature computational science* 2.10 (2022), pp. 655–664 (cit. on pp. 14, 16, 37).

[66] Hussein Mahfouz, Robin Lovelace, and Elsa Arcaute. "A road segment prioritization approach for cycling infrastructure." In: *Journal of transport geography* 113 (2023), p. 103715 (cit. on pp. 14, 16, 86, 105).

[67] Pietro Folco et al. "Data-driven micromobility network planning for demand and safety." In: *Environment and planning B: Urban analytics and city science* 50.8 (2023), pp. 2087–2102 (cit. on pp. 14, 16, 86, 105).

[68] Sheng Liu, Zuo-Jun Max Shen, and Xiang Ji. "Urban bike lane planning with bike trajectories: Models, algorithms, and a real-world case study." In: *Manufacturing & Service Operations Management* 24.5 (2022), pp. 2500–2515 (cit. on pp. 14, 16).

[69] Stefano Ardizzoni et al. "Identification of cyclists' route choice criteria." In: *2024 IEEE 63rd Conference on Decision and Control (CDC)*. IEEE. 2024, pp. 6975–6980 (cit. on pp. 14, 16).

[70] Lukas Ballo et al. "The E-Bike City as a radical shift toward zero-emission transport: Sustainable? Equitable? Desirable?" In: *Journal of Transport Geography* 111 (2023), p. 103663 (cit. on pp. 14, 16).

[71] Peter Cox. "Theorising infrastructure: a politics of spaces and edges." In: *The Politics of Cycling Infrastructure*. Policy Press, 2020, pp. 15–34 (cit. on pp. 17, 106).

[72] Peter Cox and Till Koglin. "The politics of cycling infrastructure." In: *Policy, Bristol* (2020) (cit. on pp. 17, 106).

[73] Ralph Buehler and Jennifer Dill. "Bikeway networks: a review of effects on cycling." In: *Transport reviews* 36.1 (2016), pp. 9–27 (cit. on pp. 17, 106, 109).

[74] Carole Adam and Benoit Gaudou. "An agent-based model of modal choice with perception biases and habits." In: *arXiv preprint arXiv:2406.02063* (2024) (cit. on pp. 17, 106, 109).

[75] Mehrdad Bagheri et al. "A Computational Framework for Revealing Competitive Travel Times with Low-Carbon Modes Based on Smartphone Data Collection." In: *Journal of Advanced Transportation* 2020.1 (2020), p. 4693750 (cit. on pp. 17, 106).

[76] Mehrdad Bagheri et al. "Analysis of potential shift to low-carbon urban travel modes: a computational framework based on high-resolution smartphone data." In: *Sustainability* 12.15 (2020), p. 5901 (cit. on pp. 17, 106).

[77] Andres Monzon and Luis Alfredo Vega. "How to measure the potential of transferring trips from car to public transport and soft modes—The case of Madrid." In: *Proceedings of the European Transport Conference, Strasbourg, France*. 2006, pp. 18–20 (cit. on pp. 17, 106).

[78] Andres Monzon, Luis Alfredo Vega, and Maria Eugenia Lopez-Lambas. "Potential to attract drivers out of their cars in dense urban areas." In: *European Transport Research Review* 3.3 (2011), pp. 129–137 (cit. on pp. 17, 106).

[79] John Stuart Mill. "Utilitarianism." In: *Seven masterpieces of philosophy*. Routledge, 2016, pp. 329–375 (cit. on pp. 18, 106).

[80] Ken Binmore. "Egalitarianism versus utilitarianism." In: *Utilitas* 10.3 (1998), pp. 353–367 (cit. on pp. 18, 106).

[81] John Rawls. "A theory of justice." In: *Applied ethics*. Routledge, 2017, pp. 21–29 (cit. on pp. 18, 106).

[82] Julianno Amorim and João de Abreu e Silva. "Comparing the application of different justice theories in equity analysis of transit projects: A case study of the Lisbon Metro Circular Line." In: *Journal of Transport and Land Use* 17.1 (2024), pp. 21–40 (cit. on pp. 18, 106).

[83] Syndicat Mixte des Mobilités de l'Aire Grenobloise. *Enquêtes Ménages Déplacements (EMD): Ces enquêtes sont destinées à renseigner les pouvoirs publics sur les pratiques de déplacements des habitants. Elles sont réalisées selon les méthodologies dites « standard Certu »*. Accessed: 2024-10-25. 2024 (cit. on pp. 29, 73).

[84] Crispin HV Cooper. "Using spatial network analysis to model pedal cycle flows, risk and mode choice." In: *Journal of transport geography* 58 (2017), pp. 157–165 (cit. on pp. 29, 73, 74, 78).

[85] Johannes Mesimäki and Juha Luoma. "Near accidents and collisions between pedestrians and cyclists." In: *European transport research review* 13.1 (2021), p. 38 (cit. on pp. 30, 74).

[86] Sven Maerivoet and Bart De Moor. "Transportation planning and traffic flow models." In: *arXiv preprint physics/0507127* (2005) (cit. on pp. 30, 89).

[87] United States. Bureau of Public Roads. *Traffic assignment manual for application with a large, high speed computer*. Vol. 37. US Department of Commerce, Bureau of Public Roads, Office of Planning, 1964 (cit. on pp. 31, 89).

[88] Katie Crist et al. "Fear factor: level of traffic stress and GPS assessed cycling routes." In: *Journal of Transportation Technologies* 9.01 (2019), p. 88457 (cit. on p. 37).

[89] Yang Bian et al. "Categorizing bicycling environment quality based on mobile sensor data and bicycle flow data." In: *Sustainability* 13.8 (2021), p. 4085 (cit. on p. 37).

[90] Tiziana Campisi et al. "Cycling master plans in Italy: The I-BIM feasibility tool for cost and safety assessments." In: *Sustainability* 12.11 (2020), p. 4723 (cit. on p. 37).

[91] Bettina Schröter et al. "Guidance and practice in planning cycling facilities in Europe—An overview." In: *Sustainability* 13.17 (2021), p. 9560 (cit. on p. 37).

[92] Sergio Nesmachnow and Silvina Hipogrosso. "Assessment of Sustainable Mobility Initiatives Developed in Montevideo, Uruguay." In: *Urban Science* 8.2 (2024), p. 52 (cit. on p. 37).

[93] Jisoon Lim et al. "The bicycle network improvement problem." In: *Journal of Transportation Engineering, Part A: Systems* 148.11 (2022), p. 04022095 (cit. on p. 37).

[94] CS Shui and WL Chan. "Optimization of a bikeway network with selective nodes." In: *Sustainability* 11.23 (2019), p. 6531 (cit. on pp. 37, 86, 105).

[95] Bruno Scalzo Dees et al. "Graph theory for metro traffic modelling." In: *2021 international joint conference on neural networks (IJCNN)*. IEEE. 2021, pp. 1–5 (cit. on pp. 37, 39).

[96] Samar Elaraby and Sherif M Abuelenin. "Connectivity analysis of directed highway vehicular ad hoc networks using graph theory." In: *International Journal of Communication Systems* 34.5 (2021), e4745 (cit. on pp. 37, 39).

[97] Armando Di Nardo et al. "Applications of graph spectral techniques to water distribution network management." In: *Water* 10.1 (2018), p. 45 (cit. on pp. 37, 39).

[98] Xiaofan Yang et al. "The rationality of four metrics of network robustness: A viewpoint of robust growth of generalized meshes." In: *PloS one* 11.8 (2016), e0161077 (cit. on pp. 37, 39).

[99] Ali Sydney et al. "Characterising the robustness of complex networks." In: *International Journal of Internet Technology and Secured Transactions* 2.3-4 (2010), pp. 291–320 (cit. on pp. 37, 39).

[100] Tanya Araújo and R Vilela Mendes. "Function and form in networks of interacting agents." In: *arXiv preprint nlin/0009018* (2000) (cit. on pp. 37, 39).

[101] Mohamed Bayoumi Kamel and Tarek Sayed. "The impact of bike network indicators on bike kilometers traveled and bike safety: A network theory approach." In: *Environment and Planning B: Urban Analytics and City Science* 48.7 (2021), pp. 2055–2072 (cit. on pp. 37, 39).

[102] Meisam Akbarzadeh, Syed Sina Mohri, and Ehsan Yazdian. "Designing bike networks using the concept of network clusters." In: *Applied network science* 3 (2018), pp. 1–21 (cit. on pp. 37, 39, 133).

[103] *Le Baromètre des Villes Cyclables*. `https://barometre.parlons-velo.fr/`. Accessed on January 17, 2024 (cit. on pp. 38, 149, 154).

[104] Richard H Browne. "The t-test p value and its relationship to the effect size and P (X> Y)." In: *The American Statistician* 64.1 (2010), pp. 30–33 (cit. on p. 50).

[105] Åke Björck. "Least squares methods." In: *Handbook of numerical analysis* 1 (1990), pp. 465–652 (cit. on p. 51).

[106] Jim Frost. *T-distribution Table*. `https://statisticsbyjim.com/hypothesis-testing/t-distribution-table/`. Accessed: 2024-10-03. 2024 (cit. on p. 51).

[107] Neha Patki, Roy Wedge, and Kalyan Veeramachaneni. "The Synthetic Data Vault." In: *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 2016, pp. 399–410 (cit. on p. 52).

[108] Nick Erickson et al. "Autogluon-tabular: Robust and accurate automl for structured data." In: *arXiv preprint arXiv:2003.06505* (2020) (cit. on p. 55).

[109] VVK Theja Putta. "Evaluating Bicycle Networks: Visualizing and Measuring Low-Stress Connectivity and Accessibility." PhD thesis. Northeastern University, 2019 (cit. on p. 57).

[110] Alice Brauns et al. *Le marché aménagements cyclables*. Ed. by Alice Brauns. Printed document. Source Cerema, Club de villes & territoires cyclables, 2024 (cit. on pp. 61, 78, 93, 118).

[111] Francesco Biscani and Dario Izzo. "A parallel global multiobjective framework for optimization: pagmo." In: *Journal of Open Source Software* 5.53 (2020), p. 2338 (cit. on p. 62).

[112] F. Nogueira. "Bayesian Optimization: Open source constrained global optimization tool for Python." In: *URL https://github. com/fmfn/BayesianOptimization* (2014) (cit. on p. 63).

[113] Artem Chakirov. "Sioux Falls." In: *The Multi-Agent Transport Simulation MATSim* (2016), pp. 385–388 (cit. on p. 64).

[114] Manuel Campero Jurado et al. "Improving Urban Cycling Safety and Comfort through Optimized Infrastructure Upgrades." In: *In 23rd European Control Conference (ECC)* (2025) (cit. on p. 73).

[115] M. J. Beckmann. *A Lagrangian Multiplier Rule in Linear Activity Analysis and Some of Its Applications*. Tech. rep. Paper No. 2054. Cowles Commission for Research in Economics, 1952 (cit. on p. 89).

[116] Martin Beckmann, Charles B McGuire, and Christopher B Winsten. "Studies in the Economics of Transportation." In: *The Economic Journal* 67.265 (Mar. 1957), pp. 116–118. eprint: `https://academic.oup.com/ej/article-pdf/67/265/116/27347704/ej0116.pdf` (cit. on p. 90).

[117] Pivithuru Kalpana Nanayakkara et al. "Do safe bike lanes really slow down cars? A simulation-based approach to investigate the effect of retrofitting safe cycling lanes on vehicular traffic." In: *International journal of environmental research and public health* 19.7 (2022), p. 3818 (cit. on pp. 90, 107).

[118] Marguerite Frank, Philip Wolfe, et al. "An algorithm for quadratic programming." In: *Naval research logistics quarterly* 3.1-2 (1956), pp. 95–110 (cit. on p. 93).

[119] Pramesh Kumar. *prameshk/Traffic-Assignment: Static Traffic Assignment using User Equilibrium and Stochastic User Equilibrium- Python Code*. Version 2.0. Oct. 2019 (cit. on p. 93).

[120] Rafael HM Pereira, Tim Schwanen, and David Banister. "Distributive justice and equity in transportation." In: *Transport reviews* 37.2 (2017), pp. 170–191 (cit. on p. 106).

[121] John Rawls. *A Theory of Justice*. Cambridge, MA: Harvard University Press, 1971 (cit. on p. 106).

[122] Karel Martens, Matan E Singer, and Aviv Lee Cohen-Zada. "Equity in accessibility: Moving from disparity to insufficiency analyses." In: *Journal of the American Planning Association* 88.4 (2022), pp. 479–494 (cit. on pp. 106, 115).

[123] Manuel Campero Jurado et al. "On Transport Justice and Safety in Bicycle Network Design Optimization." In: *CDC 2025-IEEE 64th Conference on Decision and Control*. IEEE. 2025, pp. 1–6 (cit. on p. 106).

[124] John Parkin. "Designing for cycle traffic." In: *International principles and practice. London: ICE Institution of Civil Engineers Publishing* (2018) (cit. on p. 108).

[125] Xingchen Yan et al. "Recommended widths for separated bicycle lanes considering abreast riding and overtaking." In: *Sustainability* 10.9 (2018), p. 3127 (cit. on p. 108).

[126] *Aménagements provisoires pour les piétons : tester pour aménager durablement*. Cité des mobilités - 25 avenue François Mitterrand - CS 92803 - F-69674 Bron Cedex, Tél. +33 (0)4 72 14 30 30. CEREMA. Bron, France, 2020 (cit. on p. 108).

[127] Geoff Boeing. "OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks." In: *Computers, Environment and Urban Systems* 65 (2017), pp. 126–139 (cit. on p. 108).

[128] Pascal Oger and Sandra Bouvet. "Aller travailler à vélo, une pratique très urbaine." In: *Insee flash Auvergne-Rhône-Alpes* (2017). ©Insee 2017 (cit. on p. 109).

[129] Agence d'urbanisme de la région grenobloise. *Enquête Mobilité 2019-2020: Chiffres clés*. Technical Report. SMMAG and CEREMA, 2020 (cit. on p. 109).

[130] Cycling Industries Europe and Megabits. *The Potential for Bike Sharing in Europe: Benchmarking 148 Cities*. Project Report. Available at: `https://cyclingindustries.com/`. Shared Ambition, Interreg North Sea Region, 2024 (cit. on p. 109).

[131] Clara Glachant and Frauke Behrendt. "Negotiating the bicycle path: A study of moped user stereotypes and behaviours in the Netherlands." In: *Transportation Research Part F: Traffic Psychology and Behaviour* 107 (2024), pp. 301–320 (cit. on pp. 109, 129).

[132] Avi Parsha and Karel Martens. "Social identity and cycling among women: The case of Tel-Aviv-Jaffa." In: *Transportation research part F: traffic psychology and behaviour* 89 (2022), pp. 1–15 (cit. on p. 109).

[133] Jeffrey M Casello and Vladimir Usyukov. "Modeling cyclists' route choice based on GPS data." In: *Transportation Research Record* 2430.1 (2014), pp. 155–161 (cit. on p. 115).

[134] Karel Martens. *Transport justice: Designing fair transportation systems*. Routledge, 2016 (cit. on p. 115).

[135] Adam Paszke et al. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." In: *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035 (cit. on p. 118).

[136] Grenoble-Alpes-Métropole. *Budget 2025: L'investissement renforcé malgré l'étau budgétaire du gouvernement*. Dossier de presse. Contact presse: Randa Berbouche, randa.berbouche@grenoblealpesmetropole.fr, 06 82 85 02 08. Mar. 28, 2025 (cit. on p. 120).

[137] Raktim Mitra, Raymond A Ziemba, and Paul M Hess. "Mode substitution effect of urban cycle tracks: Case study of a downtown street in Toronto, Canada." In: *International Journal of Sustainable Transportation* 11.4 (2017), pp. 248–256 (cit. on p. 128).

[138] Rosemary Sharples. "Travel competence: Empowering travellers." In: *Transportation research part F: traffic psychology and behaviour* 44 (2017), pp. 63–75 (cit. on p. 128).

[139] NJA Van Exel and Piet Rietveld. "Could you also have made this trip by another mode? An investigation of perceived travel possibilities of car and train travellers on the main travel corridors to the city of Amsterdam, The Netherlands." In: *Transportation Research Part A: Policy and Practice* 43.4 (2009), pp. 374–385 (cit. on p. 128).

[140] John A Hartigan and Manchek A Wong. "Algorithm AS 136: A k-means clustering algorithm." In: *Journal of the royal statistical society. series c (applied statistics)* 28.1 (1979), pp. 100–108 (cit. on p. 133).

[141] Michelle Girvan and Mark EJ Newman. "Community structure in social and biological networks." In: *Proceedings of the national academy of sciences* 99.12 (2002), pp. 7821–7826 (cit. on p. 133).

[142] Aaron Clauset, Mark EJ Newman, and Cristopher Moore. "Finding community structure in very large networks." In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 70.6 (2004), p. 066111 (cit. on p. 133).

[143] Mark EJ Newman. "Finding community structure in networks using the eigenvectors of matrices." In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 74.3 (2006), p. 036104 (cit. on p. 133).

[144] Fang Wu and Bernardo A Huberman. "Finding communities in linear time: a physics approach." In: *The European Physical Journal B* 38.2 (2004), pp. 331–338 (cit. on p. 133).

[145] Jingxin Xia and Mei Chen. "Defining traffic flow phases using intelligent transportation systems-generated data." In: *Journal of Intelligent Transportation Systems* 11.1 (2007), pp. 15–24 (cit. on p. 133).

[146] Yiming Gu et al. "An unsupervised learning approach for analyzing traffic impacts under arterial road closures: Case study of East Liberty in Pittsburgh." In: *Journal of Transportation Engineering* 142.9 (2016), p. 04016033 (cit. on p. 133).

[147] Ivana Gace et al. "Driving style categorisation based on unsupervised learning: a step towards sustainable transportation." In: *2021 6th International Conference on Smart and Sustainable Technologies (SpliTech)*. IEEE. 2021, pp. 1–6 (cit. on p. 133).

[148] Seyed Arman Haghbayan, Nikolas Geroliminis, and Meisam Akbarzadeh. "Community detection in large scale congested urban road networks." In: *PLoS one* 16.11 (2021), e0260201 (cit. on p. 133).

[149] Gary Chartrand and Ping Zhang. *A first course in graph theory*. Courier Corporation, 2013 (cit. on p. 133).

[150] David J Hand. "Principles of data mining." In: *Drug safety* 30.7 (2007), pp. 621–622 (cit. on p. 134).

[151] Ömer Nebil Yaveroğlu et al. "Revealing the hidden language of complex networks." In: *Scientific reports* 4.1 (2014), p. 4547 (cit. on p. 134).

[152] Fazle E Faisal et al. "GRAFENE: Graphlet-based alignment-free network approach integrates 3D structural and sequence (residue order) data to improve protein structural comparison." In: *Scientific reports* 7.1 (2017), p. 14890 (cit. on p. 134).

[153] Richard C Wilson and Ping Zhu. "A study of graph spectra for comparing graphs and trees." In: *Pattern Recognition* 41.9 (2008), pp. 2833–2841 (cit. on p. 134).

[154] Qun Liu, Zhishan Dong, and En Wang. "Cut based method for comparing complex networks." In: *Scientific reports* 8.1 (2018), p. 5134 (cit. on p. 134).

[155] Hisashi Kashima, Koji Tsuda, and Akihiro Inokuchi. "Marginalized kernels between labeled graphs." In: *Proceedings of the 20th international conference on machine learning (ICML-03)*. 2003, pp. 321–328 (cit. on p. 134).

[156] S Vichy N Vishwanathan et al. "Graph kernels." In: *The Journal of Machine Learning Research* 11 (2010), pp. 1201–1242 (cit. on p. 134).

[157] Cédric Villani et al. *Optimal transport: old and new*. Vol. 338. Springer, 2008 (cit. on p. 134).

[158] Cédric Villani. *Topics in optimal transportation*. Vol. 58. American Mathematical Soc., 2021 (cit. on p. 134).

[159] Filippo Santambrogio. "Optimal transport for applied mathematicians." In: (2015) (cit. on p. 134).

[160] Christoph Steinacker et al. "Robust design of bicycle infrastructure networks." In: *Scientific Reports* 15.1 (2025), p. 15471 (cit. on p. 153).

[161] R Jayakrishnan, Wei K Tsai, and Anthony Chen. "A dynamic traffic assignment model with traffic-flow relationships." In: *Transportation Research Part C: Emerging Technologies* 3.1 (1995), pp. 51–72 (cit. on p. 154).

[162] Seungkyu Ryu et al. "Two-stage bicycle traffic assignment model." In: *Journal of Transportation Engineering, Part A: Systems* 144.2 (2018), p. 04017079 (cit. on p. 154).

[163] Mads Paulsen, Thomas Kjær Rasmussen, and Otto Anker Nielsen. "Including right-of-way in a joint large-scale agent-based dynamic traffic assignment model for cars and bicycles." In: *Networks and Spatial Economics* 22.4 (2022), pp. 915–957 (cit. on p. 154).

[164] Thomas H Cormen et al. *Introduction to algorithms*. MIT press, 2022 (cit. on p. 160).

[165] Michael Barbehenn. "A note on the complexity of Dijkstra's algorithm for graphs with weighted vertices." In: *IEEE transactions on computers* 47.2 (1998), p. 263 (cit. on p. 160).