

Microsimulation and Analytical Methods to Understand Urban Air Mobility

by

Pavan Yedavalli

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

City and Regional Planning

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Paul Waddell, Chair

Professor David Culler

Professor Raja Sengupta

Professor Daniel Rodriguez

Professor Marta Gonzalez

Summer 2021

Microsimulation and Analytical Methods to Understand Urban Air Mobility

Copyright 2021  
by  
Pavan Yedavalli

## Abstract

## Microsimulation and Analytical Methods to Understand Urban Air Mobility

by

Pavan Yedavalli

Doctor of Philosophy in City and Regional Planning

University of California, Berkeley

Professor Paul Waddell, Chair

Over the last decade, a shift in mobility has occurred to accommodate different scales and purposes of travel, with accessibility becoming a more prominent objective. Micromobility, such as scooters and e-bikes, are beginning to redefine local travel. Transportation network companies, such as Uber and Lyft, are questioning the need for private vehicle ownership. Bus rapid transit and magnetic levitation are transforming the efficiency of public transit. Autonomous vehicles are reimagining the commute to accommodate medium-distance trips without the stress and lost time of driving. With technological advances in computing, machine learning, batteries, semiconductors, connectivity, and electric propulsion, urban air mobility is also becoming more popular, as it seeks to decrease the travel times for long-distance trips in sprawled and congested metropolitan areas. In short, the new key to mobility is multimodality. This dissertation focuses on the impact of urban air mobility on metropolitan transportation networks.

Introducing a new mobility mode into a region has enormous effects across a variety of metrics, including congestion, equity, sustainability, economic productivity, and long-term travel behavior and land use decisions. Since many of these objectives are often intertwined and complex, it is necessary to create tools to support infrastructure and policy decisions. This dissertation unpacks the impact of urban air mobility on travel behavior and congestion specifically, developing the foundational software simulation architecture for future scenario planning and transportation analysis.

Microsimulation Analysis for Network Traffic Assignment (MANTA), a parallelized, GPU-based microsimulation platform, is presented as an agile, modular, and extensible contribution to the transportation simulation and modeling literature. An improvement to MANTA's accuracy, namely traffic control inference using deep learning, is then proposed. Initial modeling and analytics of urban air mobility are then integrated into and output from an enhanced MANTA, and a holistic modeling framework for ground traffic, ground-air coordination, and aerial flight, known as SimUAM, is then presented. The dissertation concludes with boots

on the ground, exploring which public agencies in the United States are prepared in planning and policy to integrate urban air mobility into their transportation networks.

To Sasha, my loving wife, for supporting me through thick and thin, through uncertainty and dejection, through happiness and thrills; for not falling asleep when I describe to you what I do; for falling asleep when I describe to you what I do; for always providing the light of optimism and enthusiasm; and most importantly, for bringing our star and protector, Tarak, into the world.

# Contents

<b>Contents</b>	<b>ii</b>
<b>List of Figures</b>	<b>iv</b>
<b>List of Tables</b>	<b>ix</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.2 Context of the Study . . . . .	2
1.3 UAM’s Trajectory . . . . .	3
1.4 Organization and Contribution by Chapter . . . . .	4
<b>2 A Review of Transportation and Traffic Simulation and Modeling</b>	<b>6</b>
2.1 Introduction . . . . .	6
2.2 Background . . . . .	14
2.3 Components . . . . .	15
2.4 OD and Routing Results . . . . .	20
2.5 Simulation Results . . . . .	21
2.6 Performance benchmarks . . . . .	31
2.7 Limitations . . . . .	35
2.8 Conclusions . . . . .	35
<b>3 Machine Learning and Emergent Mobility: Advances in Transportation Simulation and Modeling</b>	<b>36</b>
3.1 Machine Learning in Simulation . . . . .	36
3.2 An Example: Traffic Control Inference using Deep Learning . . . . .	43
3.3 Autonomous Vehicles (AV) . . . . .	56
3.4 Urban Air Mobility (UAM) . . . . .	57
<b>4 Planning and Designing Land Use-Constrained Networks of Urban Air Mobility Infrastructure</b>	<b>65</b>
4.1 Introduction . . . . .	65
4.2 Methodology . . . . .	68

4.3	Modeling Results . . . . .	72
4.4	Conclusion . . . . .	79
<b>5</b>	<b>Assessing the Value of UAM through Metropolitan-Scale Microsimulation</b>	<b>82</b>
5.1	Introduction . . . . .	82
5.2	Background . . . . .	82
5.3	Methodology . . . . .	83
5.4	Results . . . . .	90
5.5	Limitations and Future Research . . . . .	92
5.6	Conclusion . . . . .	96
<b>6</b>	<b>SimUAM: A Regional-Scale Multimodal UAM Toolchain</b>	<b>97</b>
6.1	Introduction . . . . .	97
6.2	MANTA: Microsimulation Analysis for Network Traffic Assignment . . . . .	98
6.3	VertiSim: Vertiport Ground-Air Interface Modeling . . . . .	103
6.4	<i>FE</i> <sup>3</sup> : High-fidelity Aerial Simulation . . . . .	106
6.5	Simulation . . . . .	108
6.6	Performance Metrics . . . . .	109
6.7	Limitations and Future Research . . . . .	113
6.8	Conclusion . . . . .	113
<b>7</b>	<b>Evaluating the Risks and Readiness of U.S. Cities for UAM</b>	<b>115</b>
7.1	Introduction . . . . .	115
7.2	Survey Methodology . . . . .	116
7.3	Survey Findings . . . . .	118
7.4	Limitations and Additional Research . . . . .	127
7.5	Conclusion . . . . .	128
<b>8</b>	<b>Conclusion</b>	<b>129</b>
8.1	Summary of Contributions . . . . .	129
8.2	The (Sky) Ahead... . . . .	130
8.3	Final Thoughts . . . . .	132
	<b>Bibliography</b>	<b>133</b>

# List of Figures

2.1	The fundamental diagram of traffic flow that relates density, flow, and speed (courtesy of [199]) . . . . .	10
2.2	The Bay Area Network with 224,223 nodes and 549,008 edges . . . . .	16
2.3	Departure times are chosen between 5 AM to 12 PM to model the morning hours. It follows a Gaussian distribution in which the bulk of the trips begin in the 8 AM to 9 AM hour. . . . .	18
2.4	Histogram (log y-axis) showing the number of edges that see a particular vehicle count across the time range simulated. The Bay Bridge is the most heavily traversed link in the Bay Area, as it connects two major economic hubs: San Francisco and Oakland. The subset of edges representing the Bay Bridge and its necessary feeder edges (7 in total) sees a significant traffic volume of 30K trips, which has approximately 7.5K vehicles more than the next highest edges. Most edges see fewer than 100 vehicles in the timeframe. . . . .	21
2.5	A network showing all possible routes through the Bay Bridge . . . . .	22
2.6	The calibration process: average mean difference between Uber and MANTA speeds over time . . . . .	23
2.7	Comparison of trip lengths in MANTA versus California Household Travel Survey data. Median distance in MANTA is 6.46 km and in CHTS is 5.33 km. . . . .	24
2.8	5 AM - 6 AM . . . . .	26
2.9	8 AM - 9 AM . . . . .	26
2.10	Kernel density plot comparing the MANTA and Uber distributions at 35 mph . . . . .	26
2.11	Average MANTA speeds across all speed limits [5 AM - 6 AM]. The means and standard deviations are shown in parentheses. . . . .	27
2.12	Average Uber speeds across all speed limits [5 AM - 6 AM]. The means and standard deviations are shown in parentheses. . . . .	27
2.13	Average MANTA speeds across all speed limits [8 AM - 9 AM]. The means and standard deviations are shown in parentheses. . . . .	28
2.14	Average Uber speeds across all speed limits [8 AM - 9 AM]. The means and standard deviations are shown in parentheses. . . . .	28
2.15	Average MANTA speeds across all speed limits [8 AM - 9 AM] in the red light case. The means and standard deviations are shown in parentheses. . . . .	29
2.16	Fit to lognormal distribution for 20 mph speed limit in green light scenario (case 2)	30



2.17	Fit to normal distribution for 45 mph speed limit in green light scenario (case 2)	30
2.18	Time required to route agents using priority-queue Dijkstra algorithm for the SF Bay Area network on distributed computing environment (MPI + OpenMP) parallelization. Tests were run on 32 nodes with Intel Xeon Skylake 6142 processors.	32
2.19	Speedup of priority-queue routing algorithm for the Bay-Area network on distributed computing environment (MPI + OpenMP) parallelization. Tests were run on 32 nodes with Intel Xeon Skylake 6142 processors.	33
2.20	Simulator runtimes (log scale y-axis) across different simulators. MANTA performs slightly better than the parallelized mesoscopic JDEQSIM and is on the same order of magnitude. MANTA performs significantly better than the mesoscopic version of SUMO with either the simplified (MeS) or advanced intersection modeling (MeA). The microscopic version of SUMO with simplified intersections (MiS) and advanced intersections (MiA) could not be run completely, and thus times were linearly extrapolated, reflecting that it would take tens of days to complete.	34
3.1	All nodes in the bounding box of the Nokia HERE data	44
3.2	All nodes with traffic control in the bounding box of the Nokia HERE data	45
3.3	Sample trajectory in space from a particular device	45
3.4	Sample node with its associated trajectory points across time, distance from node, and speed.	47
3.5	Sample node with its associated trajectory points across time, distance from node, and speed, linearly interpolated to create a heatmap.	47
3.6	Training accuracy using ReLU as the activation function for all layers except the output layer	49
3.7	Training loss using ReLU as the activation function for all layers except the output layer	50
3.8	Training accuracy with 50, 100, and 200 filters for the convolutional layers (with ReLU activation function)	51
3.9	Training loss with 50, 100, and 200 filters for the convolutional layers (with ReLU activation function)	51
3.10	A random set of 30 test images shows success rate of approximately 70% overall with ReLU activation.	52
3.11	Training accuracy with 50, 100, and 200 filters for the convolutional layers (with tanh activation function)	53
3.12	Training loss with 50, 100, and 200 filters for the convolutional layers (with tanh activation function)	53
3.13	A random set of 30 test images shows success rate of approximately 72% overall with tanh activation.	54
4.1	The San Francisco Bay Area Network	68

4.2	In San Francisco, a sample vertiport parcel (red) determined from the clustering algorithm and its nearest node (orange) that is used in subsequent microsimulation	72
4.3	Converged map of $k = 10$ original cluster centroids across the San Francisco Bay Area, in radians, based on OD nodes of the street network . . . . .	73
4.4	Distances between calculated clusters and feasible parcels shows that as $k$ increases, the mean distance tends to decrease. The green triangle represents the mean distance and the orange line represents the median distance. The box and edges show interquartile range, minimum, and maximum distances. . . . .	74
4.5	Adjusted cluster centroids (vertiport parcels) across the San Francisco Bay Area that satisfy local land use requirements for $k = 10$ (blue) and $k = 50$ (orange) .	75
4.6	Difference in travel times between driving and UAM across $k$ when parcels constrained to industry, retail, office, and parking garages, and to 2 acres or larger. Positive values imply longer non-UAM travel times and negative values imply longer UAM travel times . . . . .	76
4.7	Mean travel times across $k$ in UAM scenarios vs. mean travel time in a non-UAM scenario when parcels constrained to industry, retail, office, and parking garages and to 2 acres or larger. For $10 \leq k \leq 30$ , UAM takes longer than driving for the same trip. For $40 \leq k < 100$ , driving takes longer than UAM for the same trip. $k \geq 100$ produces shorter mean travel times across all trips. . . . .	77
4.8	The origin, destination, and vertiports in an example UAM trip. Origin in peach, destination in green, vertiport nearest the origin in red, and vertiport nearest the destination in light blue. This trip would not benefit from UAM because the vertiports are far from the origin and destination and in opposite directions of the desired direction of travel. The driving distance between origin and destination is 15 miles, while the driving and aerial distances together with UAM is approximately 38 miles. . . . .	77
4.9	The origin, destination, and vertiports in an example UAM trip. Origin in blue, destination in orange, vertiport nearest the origin in brown, and vertiport nearest the destination in magenta. This trip benefits from UAM because the vertiports are along the way if the trip was made driving. The driving distance between origin and destination is 18.6 miles, while the driving and aerial distances together with UAM is approximately 15.5 miles. . . . .	78
4.10	Mean travel times for UAM and non-UAM scenarios as area constraints of vertiport parcels varies. When $A = 2$ acres, mean UAM travel times are lower than non-UAM, but as $A > 2$ , mean travel times are at least 20 minutes higher for UAM than for all trips while driving. . . . .	78
5.1	A typical UAM flight trajectory, adapted from [191] . . . . .	84
5.2	The travel analysis zones (TAZ) in the San Francisco Bay Area . . . . .	85

5.3	In (a), the histogram of trips across distance from 5AM-12PM and from 8AM-9AM. The average distance and standard deviations are similar. In (b), the number of trips relative to the inter-vertiport distance for trips between 8AM and 9AM, showing a steep decline in trips as the distance increases. . . . .	86
5.4	Network designs determined from k-medians clustering at $k = 10$ and $k = 100$ . . . . .	87
5.5	Distances . . . . .	88
5.6	Every multi-modal UAM trip contains an access leg, transfer at the origin vertiport, aerial leg, transfer at the destination vertiport, and an egress leg. . . . .	88
5.7	Iterative convergence across multiple simulations narrows down the number of trips that benefit from taking multi-modal UAM . . . . .	89
5.8	The number of trips that benefit from multi-modal UAM . . . . .	91
5.9	The number of benefited trips across $k = 10, 30, 100$ after convergence . . . . .	93
5.10	At $k \geq 10$ , the average travel time on the road is lower than driving those same trips, showing that there is benefit to having UAM in the transportation network once there are at least 10 vertiports, with significant benefits occurring at 30 vertiports or more. . . . .	94
5.11	Trips that do not take UAM also benefit, albeit slightly, when UAM is an available mode in the transportation network. . . . .	95
5.12	At $k = 10$ , there is an imbalance among vertiports in how many trips are served . . . . .	95
6.1	Simulation components for an individual multi-modal UAM trip . . . . .	98
6.2	The full SimUAM model pipeline . . . . .	98
6.3	The San Francisco Bay Area region's edges and nodes, defined by the polygonal hull of its nine counties . . . . .	100
6.4	The MANTA architecture . . . . .	102
6.5	Kernel density estimator plot comparing MANTA microsim edge speeds vs. Uber edge speeds on edges whose speed limits are 45 mph . . . . .	103
6.6	The ground/air interface at the vertiport . . . . .	105
6.7	FE <sup>3</sup> aerial microsimulator architecture . . . . .	106
6.8	The number of benefited trips when the ground-air interface is not modeled . . . . .	109
6.10	The number of benefited trips once the ground-air interface is simulated in SimUAM significantly decreases across all $k$ . . . . .	110
6.11	The average travel time of benefited trips once the ground-air interface is simulated in SimUAM significantly increases across all $k$ (red vs. orange) . . . . .	111
6.12	The average transfer time at $k = 10$ once the ground-air interface is considered . . . . .	111
6.13	The average transfer time and average aircraft occupancy with different demand levels . . . . .	112
6.14	The congested airspace of the San Francisco Bay Area at $k = 30$ vertiports . . . . .	113
7.1	Respondents by jurisdiction type and by role within their organization . . . . .	117
7.2	Respondents by jurisdiction type and location across the United States . . . . .	118

7.3	The Relationship between local communities, aviation stakeholders, and UAM community integration . . . . .	119
7.4	Average priority level of equity concerns across all jurisdiction types. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned. . .	123
7.5	Average priority level of integrating with existing modes across all jurisdiction types. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned. . . . .	124
7.6	Average priority level of safety concerns across cities, states, and MPOs. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned.	125
7.7	Average priority level of sustainability concerns across cities, states, and MPOs. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned. . . . .	126
8.1	Incorporating the long-term land use loop . . . . .	131

# List of Tables

2.1	IDM parameter ranges, derived from [213] . . . . .	19
2.2	MANTA's runtimes compared to SUMO and JDEQSIM. Full implies that the entire simulation was able to complete. Lin. extrap. implies that only part of the simulation was able to complete and the full time was linearly extrapolated from this preliminary time. . . . .	34
4.1	Taxonomy and Definitions of UAM Infrastructure (adapted from Cohen et al.) .	66
7.1	UAM Planning by jurisdiction type (Notes: (1) Percentages per column will not add up to 100% due to how the data is presented. Instead, the sum of percentages in each column will signify the total percentage of respondents who answered Yes to the question, with the remaining percentage as No or Unsure), (2) Each percentage is rounded to the nearest whole number. . . . .	121

## Acknowledgments

A Ph.D, in many respects, is an irrational decision, a labor of love in the pursuit and creation of new knowledge. While the journey at times may feel solitary, the Ph.D is always a group effort, and I was fortunate enough to have advisors, mentors, colleagues, friends, and family constantly cheering me on.

I'd like to first and foremost thank my primary advisor, Dr. Paul Waddell. He took a blind chance on me four years ago after I e-mailed him showing interest in his research. He saw my research potential, and my passion and excitement for incorporating technology responsibly into the urban ecosystem. He then gave me amazing academic conversations and insights, and the freedom to find my interests. From the beginning, he has nurtured my entrepreneurial spirit. He is a role model, a mentor, and a friend at the same time, a combination that is rare to find in any individual. I am forever grateful for his unwavering guidance, support, and inspiration. I'd also like to give a special thank you to Dr. Raja Sengupta for being a fantastic committee member and advisor over the last year. He gave me the green light to lead the Cal Unmanned Lab research group during our work with NASA, and I was able to develop and cultivate relationships with the group that I will forever cherish. I also appreciate all the learnings and guidance from the rest of my committee: Dr. David Culler, Dr. Marta Gonzalez, and Dr. Daniel Rodriguez. I would like to give thanks to Dr. Min Xue at NASA as well. He and I started a casual lunch conversation down at Ames about how our two simulators could work nicely with each other, and ended with an exciting year-long UAM project under the auspices of NASA and Berkeley. It truly is amazing what a fortuitous meet-up can lead to.

I'd like to thank my co-authors contributing to various papers in this dissertation - Dr. Krishna Kumar, Dr. Jane Macfarlane, Emin Burak Onat, Xin Peng, Adam Cohen, and Dr. Vishwanath Bulusu. Finally, a special thanks must be given to all of my wonderful friends and colleagues at Berkeley and UrbanSim along the way - Ignacio Garcia Dorado, Max Gardner, Juan Caicedo, Julian Palladino, Dr. Sam Maurer, Chester Harvey, Dr. Abigail Cochran, Savannah Cox, Yanin Kramsky, Giselle Mendonca, Priscila Rochi, Megumi Yamanaka, Clara Turner, Theophile Cabannes, and Dr. Bingyu Zhao - they gave me support, criticism, guidance, laughs, and most importantly, demonstrated patience as I tried to describe to them that a Jetsons future may actually be around the corner.

Of course, no acknowledgement section would be complete without thanking my family. My 19-month-old son and little monster, Tarak Yedavalli, has kept me disciplined, grounded, and most importantly, sane, as I navigated the home stretch of the Ph.D. Running around with him, taking him on Saturday morning walks to a new coffee shop in San Francisco to get an espresso, kissing his little belly, and seeing him laugh made and continue to make every possible challenge worth it. I'd also like to thank my wife, Dr. Sasha Gupta, to whom this dissertation is dedicated. She has been with me through the thick and thin after we met on this very Berkeley campus fourteen years ago. Her love, positivity, and support has inspired me to increase the height of my ceiling while simultaneously keeping my floor high.

The Ph.D is not only a labor of love, but it is a test of love, and to have my better half right by my side at every point made this journey that much more fulfilling.

I'd also like to thank my parents - my father, Dr. Rama K. Yedavalli, who has been a professor of aerospace engineering at The Ohio State University for over 30 years, and my mother, Sreerama Yedavalli, who was also in the middle of doing her Ph.D in sociology before having kids. They showered me with absolutely everything I could have asked for and more, indulging my curiosities, pushing me in subtle yet inspiring ways, and providing me the comfortable and loving environment that helped me grow into the person that I am today. The intellectual atmosphere at home, the fun of going to my dad's office on campus, and the constant FBI questioning that my mom always used not only helped me understand that being thoughtful and inquisitive, respectfully, is always a good thing. I'd also like to thank my brother, Dr. Vivek Yedavalli, who, at almost 3 years older, has always been my best friend and role model, pushing me, testing me, playing with me, and always making me and my parents laugh. My parents and brother are the reasons I've believed in pursuing knowledge, and they've provided everything to make it possible. I'd also love to thank my sister-in-law, Aditi Yedavalli, and my two twin nephews, Aatrey and Vaibhav Yedavalli, whose smiles whenever their uncle FaceTimes them always make my day. And I'd finally love to thank my in-laws: father-in-law, Vivek Gupta; mother-in-law, Mamta Gupta; sister-in-law, Dr. Charu Gupta Soni; brother-in-law, Ravin Soni; and my nieces, Gia and Mina Soni, who have always been supportive of my academic journey and have provided the laughs, positivity, and warmth.

There are several technical frameworks that deserve a special thanks as well. Intelligent driver models, parallelized CUDA frameworks, and stellar documentation have made the bulk of this research possible, on both the C++ development side and Python analytics side. The bevy of open-source knowledge and information that exists in the world in order to assist software development, from StackOverflow to machine learning and data science blogs to useful Python libraries, is criminally underappreciated. I could not have learned and moved forward with my research without these resources.

Finally, this research would not have been possible without the generous financial support of the U.S. Department of Energy SMART Mobility initiative, National Aerospace and Science Administration (NASA), Universities Space Research Association (USRA), and the U.S. Department of Transportation Dwight D. Eisenhower Transportation Fellowship.

# Chapter 1

## Introduction

### 1.1 Motivation

The history of large-scale aviation begins in the early 20th century, when twin jet engines, with cowling around their circumferences for greater thrust efficiency and less drag, were developed for larger planes for use in World War I, almost fifteen years after the Wright brothers developed the first powered, manned, sustained flight aircraft [94]. Major aircraft manufacturers such as Boeing and Douglas developed large-scale aircraft with powerful turbojet engines, transporting hundreds of people per aircraft [218]. This fixed-wing aviation technology continues to prevail for long-distance travel, with a global commercial aviation industry enabling millions of people to travel across the world every day [218]. On the other hand, rotor aircraft, such as helicopter, have largely remained niche aerial vehicles due to their inefficiency against drag and when gaining lift, compounded by their relatively prohibitive cost [218]. However, over the past two decades, rotors have become more aerodynamic, and this greater efficiency has allowed for alternative propulsion techniques to fuel, such as distributed electric propulsion [58, 221]. Urban air mobility (UAM) is a product of these technical advancements, and the forthcoming electric vertical takeoff and landing (eVTOL) aircraft are expected to leverage these new innovations for passenger and cargo transport [58, 221].

Proponents believe that UAM will overcome transport inefficiencies by flying in lower altitude airspace, moving quickly from point A to point B. The current architecture for UAM envisioned by governmental agencies, academics, and industry leaders is a shared service in which consumers request an on-demand aerial taxi, similar to the experience of ridesharing services today [58, 77, 220, 194, 206]. By supplementing inter- and intra-city transportation demand with UAM, advocates anticipate that travel times will be greatly reduced, and long-run positive downstream effects may include lesser environmental degradation, decreased energy costs, and fewer resources spent on infrastructure by leveraging free space.

Indeed, this is merely speculation rooted in excitement rather than objectivity. In reality, UAM's impacts are wholly unknown, and may be positive, negative, or both, depending on



the geographic context, population investigated, and/or metric analyzed. In order to plan for UAM as a new mode, we must first understand the range of UAM's impacts on a metropolitan region and its transportation network. To do so, we must do extensive simulation before deployment, which will then inform planners, policymakers, and engineers as they conduct cost-benefit analyses. This dissertation explores analytical methods of understanding UAM's impact on metropolitan regions, discerning whether UAM can indeed be the panacea to congestion that proponents believe, and if so, in which deployments. Novel contributions in regional-scale microsimulation and traffic control inference are first presented. UAM is then integrated into this microsimulation framework, using the San Francisco (SF) Bay Area as a case study. Feasible vertiport network designs under land use constraints as well as detailed regional congestion analyses with varying aerial and vertiport assumptions are presented. The dissertation concludes with a mixed-methods study on the preparedness of local, regional, and state agencies for UAM, and the transportation planning implications for this new mode.

## 1.2 Context of the Study

With technological advances in computing, machine learning, batteries, semiconductors, and connectivity, a shift in mobility is occurring to accommodate different scales and purposes of travel, with accessibility becoming a more prominent objective. Micromobility, such as scooters and e-bikes, are beginning to redefine local travel. Transportation network companies (TNCs), such as Uber and Lyft, are questioning the need for private vehicle ownership. Bus rapid transit and magnetic levitation are transforming the efficiency of public transit. Autonomous vehicles (AV) are attempting to reimagine the commute to accommodate medium-distance trips in the comfort of an automobile without the stress and lost time of driving. UAM as a concept is also becoming more popular, as it seeks to decrease the travel times for long-distance trips in sprawled and congested metropolitan areas. In short, the new key to mobility is multimodality. For every trip purpose, every person's preferences, and every distance traveled, experts believe there should be a safe, sustainable, equitable, and efficient mode of travel available [50].

However, introducing a new mobility mode should assume a level of diligence that the mode will be beneficial to most, if not all, people. Since UAM is an early concept, safety, sustainability, equity, and efficiency cannot be guaranteed, and in fact, may become worse if UAM is deployed irresponsibly. Effective deployment begins with comprehensive transportation planning and modeling, primarily through travel demand models and simulation [123]. Travel demand modeling has often been explored through two complementary lenses: transportation models and traffic flow models [136]. Maerivet et al. discuss how transportation models operate at a high level, analyzing how households make certain travel decisions about social, economic, and cultural activities, which then inform transportation policy and infrastructure decisions [136]. There are two types of transportation model approaches: (1) trip-based, and (2) activity-based.

While transportation models explore where travelers go and why, traffic flow models describe how people are moving in the pipes themselves, specifically investigating the physical propagation of travelers through a road network. There are traditionally three types of traffic flow approaches: (1) macroscopic, (2) mesoscopic, and (3) microscopic, depending on the level of traveler aggregation [228, 210, 182].

Both transportation and traffic flow models operate at the same levels of granularity, from the coarsest aggregate macroscopic trip-based models to the most zoomed-in disaggregate microscopic agent-based models. Simulation models encompass the gamut of different analyses, often using aggregate and disaggregate analyses together [101]. When technological capabilities exist, simulation models are extremely valuable in scenario planning and impact analysis [101, 226].

Introducing a new mobility mode into a region has enormous effects across a variety of metrics, including congestion, equity, sustainability, economic productivity, and long-term travel behavior and land use decisions. While these objectives are often intertwined, it is valuable to investigate each independently to understand UAM’s upper and lower bounds. Using simulation, this dissertation attempts to unpack the impact of UAM on congestion specifically, developing the foundational software architecture and testing framework by which the subsequent measures can also be tackled in future research.

In the bulk of this work, the San Francisco (SF) Bay Area is used as a case study. This is for two reasons. The first is the unique nature of economic and spatial growth of the region, which produces interesting commuting patterns. Specifically, because the region originally developed from the gold-boom town of San Francisco, it had assumed patterns of monocentrism, in which most economic activity was centralized in the city, and thus much of the travel, on all historic modes, was toward the city in the morning and away from the city in the evening [45, 103]. As the region prospered over the twentieth and twenty-first centuries, Oakland to the east emerged as an industrial hub, and San Jose, to the south and the epicenter of Silicon Valley, emerged as a commercial and residential hub as well, producing a new polycentrism in which there are multiple centers of economic and cultural activity [45, 102, 121, 98, 134, 214]. Today, travel patterns in the SF Bay Area reflect both this monocentrism and polycentrism, which makes the region’s findings highly adaptable to either structure of any city across the world.

The second reason is functional. Researchers at University of California, Berkeley have extensive access to both public and private travel data collectors across the region, including the Bay Area Metropolitan Planning Organization, San Francisco Municipal Transit Authority, California Department of Transportation, and major mobility providers such as Google, Uber, Lyft, and Apple, all headquartered in the SF Bay Area.

### 1.3 UAM’s Trajectory

Before simulating and understanding UAM’s impacts, the challenges to deployment must be highlighted. Historically with many new technologies, deployment is incremental; UAM’s

implementation will likely follow a similar trajectory [100]. UAM is initially expected to be expensive but become more affordable over time, as eVTOL aircraft technology becomes more commoditized [100]. As traffic in the air remains low, the bottlenecks and operational issues with unmanned traffic management (UTM) overseen by air traffic control will also evolve [20].

Since UAM is in its rudimentary stages, developing the operational theory and logistical architecture itself is challenging. Many stakeholders will play major roles in initial deployment [39]. The aircraft are likely to be manufactured by leaders in industry, similar to commercial aircraft today, and operations may be managed by private entities [58, 77, 194]. Duffy et al. note that logistics and infrastructure will be developed and maintained by public agencies and businesses, but subsidized by private entities that rent out space, similar to present-day airports [75]. Song et al. predict that operation and optimization will be regulated by aviation regulatory bodies, such as the FAA and EASA [196]. Balac et al. suggest that the stops, known as vertiports, are also expected to be dispersed across metropolitan regions, initially leveraging existing helipad locations, and eventually being placed in highly connected areas where travel demand is dense, similar to subway and commuter rail [20].

There are several challenges to the adoption of UAM as well. Public perception, specifically related to safety and noise, must be managed in order to gain consumer trust and confidence. The overflow are particularly concerned about the vehicles unpredictably falling out of the sky [241]. Aircraft noise will also be a major parameter in community acceptance, and it may dictate and restrict access to vertiports in specific geographic areas or certain times of the day [220]. Different studies also show that privacy is a major concern for communities, as UAM is often associated with military drones and increased government oversight [47, 53]. In addition, political polarization can dictate public opinion about unmanned aerial vehicles (UAV), under which UAM is a subset [172].

Ultimately, UAM faces a challenging path to deployment, both in the United States and across the world. Overcoming perception issues, developing viable business models, and understanding its impacts on a variety of important transportation and societal issues are paramount, and will ultimately dictate whether UAM takes off at scale.

## 1.4 Organization and Contribution by Chapter

In this dissertation, we develop a holistic simulation toolchain that investigates UAM’s impacts on street network congestion, with the hope that the framework will become a foundation for future scenario planning and transportation analysis. The toolchain’s architecture and functionality are discussed in the subsequent chapters, with each chapter building on the previous.

Chapter 2 introduces the field of travel demand modeling and presents the prevailing paradigms, including the canonical four-step model (4SM). This model consists of trip generation, trip distribution, mode choice, and traffic assignment across household counts, classified by geography, income group, or other simple categories [96, 136, 118, 199, 12]. More

disaggregation and granularity has recently been explored in activity-based models, which are microscopic in nature and often provide a more direct representation of real-world human behavior [227, 238, 210, 212, 236]. The chapter continues by exploring traffic flow theory and the physical propagation of vehicles to understand their movements. Microsimulation Analysis for Network Traffic Assignment (MANTA), a parallelized, GPU-based microsimulation platform, is presented as an agile, modular, and extensible contribution to the transportation simulation and modeling literature.

Chapter 3 introduces new enhancements and disruptions to the aforementioned modeling and simulation paradigms. It begins with a discussion of machine learning (ML) and how it has augmented current simulation tools, detailing deep learning algorithms in the context of simulation and traffic flow prediction. A novel example of machine learning to improve simulation capabilities, specifically using convolutional neural networks to infer traffic control at intersections, is then presented. Finally, emergent mobility modes and their impact on simulation and modeling, from the perspectives of data collection and changes in travel behavior, are discussed.

Chapter 4 reintroduces UAM as a mobility paradigm and describes the integration of UAM into microsimulation. First, the planning and analysis of vertiport siting is presented. Then, using the SF Bay Area as a case study, a regional travel time analysis of UAM with a vertiport network design constrained by land use is given.

Chapter 5 provides additional analysis that incorporates a data-driven approach to determine the total addressable market of UAM, as well as analytics across transfer times and number of vertiports. It ends with a discussion about regional congestion with UAM as part of the transportation network. This chapter provides a baseline for future work in UAM mode choice modeling.

Chapter 6 replaces the initial assumptions made in the analyses of Chapter 4 with a full-scale integrated multimodal toolchain, SimUAM, that simulates all components of a UAM trip in the transportation network. This includes access and egress legs to the vertiports, transfer and consolidation at the vertiports, and aerial travel using particular schedules as well as conflict detection and resolution. By replacing static assumptions with dynamically simulated ones, bottlenecks are more readily identifiable.

Chapter 7 zooms out to understand the planning requirements of integrating UAM into the transportation network. It first introduces the challenges associated with deploying UAM in metropolitan areas, citing significant issues such as equity, safety, and sustainability, among others. A mixed-methods approach, with the development of a survey instrument as well as the subsequent analytics, is then used to evaluate the preparedness of domestic government agencies, including cities, metropolitan planning organizations (MPOs), and states, for UAM.

Chapter 8 provides a conclusion of the work presented in the dissertation, as well as future outcomes and research necessary to advance UAM simulation.

## Chapter 2

# A Review of Transportation and Traffic Simulation and Modeling

Mobility simulation has historically revolved around travel demand modeling to understand how and to where humans move on both temporal and spatial scales [226]. Travel demand modeling has often been explored through two complementary lenses: transportation models and traffic models [136]. The chapter will discuss in turn each lens.

## 2.1 Introduction

### Transportation Models

There are two types of transportation model approaches: (1) trip-based, and (2) activity-based [182]. The transportation model that largely dominates travel demand modeling is the aggregate four-step model (4SM), first formalized as a relationship between activity patterns and transportation systems by Mitchell and Rapkin, as well as Beckmann et al. in the 1950s and 1960s [136]. The 4SM first requires trip generation, in which the total number of trips to and from a particular unit of geography is assigned. The second step is trip distribution, in which the origins and destinations of these specific trips are determined [182, 123]. These first two steps are straightforward byproducts of data collection.

### Mode Choice and Choice Models

The third step is to evaluate the mode choice, selected among the set that includes the car, bike, foot, public transit, and other forms of transport. Mode choice itself is a simple artifact of the data, but it is more important to understand the parameters that affect mode choice. Mode choice is quantitatively measured by both the expected absolute number of users for each mode in a particular geography, and the probability that these users take each mode. In order to understand the mode choice distribution, multinomial logit (MNL) models are commonly used [136, 29].

MNL models required in the third step of the 4SM are based on the microeconomic theory of utility maximization, in which humans are believed to be rational actors who make decisions that maximize their total value derived from a particular decision [29, 211, 210]. While research in psychology and behavioral economics has shown evidence that human behavior is highly varied, dependent on individual motives, social and cultural norms, and even groupthink, the assumption of rationality continues to be used as a basic foundation, for the sake of simplicity and interpretability [136, 141]. Each mode's utility contains a systematic component and a stochastic component, as follows:

$$U_{drive} = V_{drive} + \epsilon_{drive} \quad (2.1)$$

where  $V_{drive}$  is the systematic and observable utility and  $\epsilon_{drive}$  is the random error component. In order for a traveler to pick driving against public transit or rideshare, for instance, we assume that their  $U_{driving} > \max(U_{PT}, U_{RS})$  where  $PT$  is public transit and  $RS$  is rideshare. The probability of person  $n$  choosing to drive over other modes is specified as follows:

$$P_n(drive|X_n) = \frac{e^{V_{drive}}}{e^{V_{drive}} + e^{V_{PT}} + e^{V_{RS}}} \quad (2.2)$$

where  $X_n$  represents the attributes of the person,  $PT$  is public transit, and  $RS$  is rideshare. As with all probabilities, multinomial logit probabilities arise from dividing that choice by the total number of choices. Determining these probabilities requires specifying utility equations for each mode [29, 211]. Since we cannot observe the error  $\epsilon$ , the only understanding we have of a person's utility for that mode is in the systematic component,  $V$ . It is important to note that, in logit models, these probabilities can be derived only under the assumption that the random error term of each alternative's utility is independent and identically distributed (iid) and follows the extreme value Gumbel distribution.  $V$  for a particular person is specified as a linear combination of weighted parameters, as follows:

$$V_{i,n} = \beta_0 + \beta_1 x_{i,n} + \beta_2 x_{i,n} + \dots + \beta_k x_{i,n} \quad (2.3)$$

where  $n$  is the specific person,  $i$  is the alternative or choice,  $x_{i,n}$  is the parameter included in the utility equation, and  $\beta_k$  is the coefficient that represents the weighting of that parameter in the person's overall decision to choose that alternative  $i$ .

In order to estimate these values, data are collected first. A stated or revealed preference survey is typically designed to capture consumer demand for a particular mode choice. The survey helps researchers understand current individual commuting habits and demographic information of the respondent, generated using an orthogonal design [29, 211]. Several sample scenarios are usually presented, with each respondent's choosing their preferred mode choice out of a given set of modes. Each alternative has several attributes, such as travel time and price, each modeled by an  $x_i$  in Equation (2.3). The survey data must then be extensively cleaned and prepared for the MNL specification above. The coefficients,  $\beta_k$ , are then calculated based on maximum likelihood estimation. The likelihood of person  $n$  choosing alternative  $i$  is specified as follows:

$$L^* = \prod_{n=1}^N P(Y_n|x_n; \beta) \quad (2.4)$$

where  $Y_n$  is the chosen alternative,  $x_n$  represents all the explanatory variables, and  $\beta$  represents all coefficients. The probability  $P(i)$  is the likelihood that person  $n$  chooses alternative  $i$ . Given this equation, we want to maximize the likelihood  $L^*$  of predicting every person's mode choice, which then indicates a better model specification for that alternative. First-order optimization tells us that in order to maximize, we must differentiate with respect to  $\beta$ . Once we are able to estimate the various  $\beta$  coefficients for each alternative, then we can assess each parameter's statistical significance by analyzing its standard error and p-value. This is often an exploratory and iterative process to determine the proper model specification for each mode choice [29, 211]. Models are then compared to one another using likelihood ratio tests and assessments of values of time (VOT), along with  $\rho^2$  values.

## Traffic Assignment

The final step in the 4SM is traffic assignment, in which the specific paths in the network from origin to destination are established. This is often done by finding a stable equilibrium using shortest-path algorithms based on such impedances as travel time, distance, or cost (from a metric such as fuel consumption) [136, 199, 204]. Since individuals' route choices affect others' choices by imposing higher travel costs in the form of slower travel, traffic assignment ideally should be modeled as a dynamic process, but since this often involves more complexity, static traffic assignment, during which impedances are calculated only once, is used. Unsurprisingly, static traffic assignment (STA) typically falls short in representing the most realistic behavior on the network [136, 199, 204].

In order to understand how travel behavior is dictating congestion, we can calculate travel times for each edge in a given region's street network by simulating travel across all transportation networks over a specified number of timesteps. This can involve aggregate and disaggregate approaches, using both static traffic assignment and dynamic traffic assignment (DTA), respectively. In STA, a map of congestion is created over time and space. Often, to relate the edge travel time and edge congestion, the standard Bureau of Public Roads (BPR) congestion function is used if data do not exist:

$$t_i = t_0 \left( 1 + \alpha \left( \frac{v_i}{c_i} \right)^\beta \right) \quad (2.5)$$

where  $t_i$  is the congested flow travel time,  $t_0$  is the free-flow travel time,  $v_i$  is the number of vehicles of edge  $i$  per unit time, and  $c_i$  is the capacity (maximum number of vehicles) of edge  $i$  per unit time [204]. The constant  $\alpha$  is a linear multiplier while  $\beta$  is an exponential multiplier of the volume-to-capacity ratio. In order to calculate  $\alpha$  and  $\beta$ , we need edge data about capacity and free-flow travel time. For simplicity, these are typically set as .15 and 4, respectively, based on canonical data. Edge lengths, free-flow speeds, and each edge's number

of lanes are often obtained from such libraries as OpenStreetMap (OSM) and OSMnx, and probe data for speeds are obtained from mobility providers. Once the road network, each edge's BPR coefficients, and node-based travel demand data are established, then the traffic is simulated over a number of time iterations in order to obtain  $t_i$ .

This macrosimulation requires a first-order optimization routine, such as the Frank-Wolfe algorithm, which allocates traffic to each edge under a certain equilibrium condition, such as user equilibrium (Wardrop's first principle), shown below; or system equilibrium (Wardrop's second principle), shown further below [112].

$$\min(Z) = \sum_a \int_0^{x_a} t_a(x) dx \quad (2.6)$$

or

$$\min(Z) = \sum_a \int_0^{x_a} x t_a(x) dx \quad (2.7)$$

subject to

$$\sum x_a = C \quad (2.8)$$

where  $t_a$  is the travel time on link  $a$ ,  $x_a$  is the total number of flows on link  $a$ , and  $C$  is a constant for total number of flows in Eq. 2.8. Colak et al. find that 15-30% of total time lost in congestion is caused solely by selfish routing [56].

In simulation, DTA edge impedances must be constantly calculated and fed back into the optimization or routing engine, often using heuristics such as time increments for calculation.

## Traffic Flow Models and Theory

The last step of traffic assignment makes way for an area of scholarship that requires its own detail: traffic flow models.

### Fundamental Diagram of Traffic Flow

Traffic theory has historically developed in the context of automobiles on roads and highways, but can also be applied more broadly across other forms of transport. This section focuses on the automobile because of its prevalence and relevance to the overarching transportation network within a metropolitan area. Traffic flows are often analyzed from a fluid mechanics perspective [12]. Three major variables of traffic exist: flow, density, and speed (or velocity) [199, 12], all represented in the fundamental theory of traffic flow, shown in Equation (2.9):

$$q = kv \quad (2.9)$$



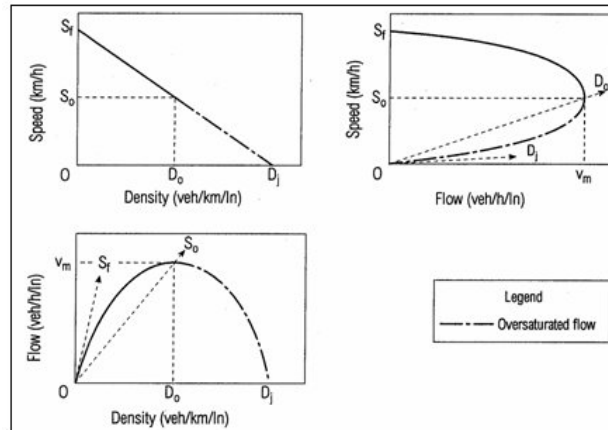


Figure 2.1: The fundamental diagram of traffic flow that relates density, flow, and speed (courtesy of [199])

where  $q$  is flow,  $k$  is density,  $v$  is speed (velocity). Flow is measured as the quantity of traffic that moves from one point to another over a period of time. More specifically, it can be the number of people, goods, or vehicles passing a point in a set period of time [199, 12]. Density is measured as the concentration of vehicles per unit distance per unit time. When density is low, then the count of vehicles along the edge is small. When density is high, that implies that vehicles are very close to one another and hardly moving [199, 12]. If they are hardly moving or come to a full stop, the flow goes to zero, which means that the density is at a maximum while the flow is zero. Speed can be classified as instantaneous speed or mean speed [199, 12]. The instantaneous speed is given by the slope at a given point in the space-time diagram, while the mean speed, which is more commonly used, occurs when averaging the speed of all vehicles over a period of time at a particular point on the transport network [199, 12]. When density is low and the flow is low, then the speed will likely be at a maximum, representing free-flow speed. As flow increases, then the density also increases, which then forces speeds to be adjusted to maintain safe distance headways. As the density increases, speeds continue to drop; when jam density is reached, the speed halts to zero. Similarly, as flow increases, speeds will decrease, but generally more slowly over a large range of flows [199].

The first relationship, known as the fundamental diagram of traffic flow, shows density on the x-axis and flow on the y-axis, and is a concave down graph in which the flow and density increase until a maximum, shown in Figure 2.1, at which point the flow decreases, despite density continuing to increase. The inflection point corresponds to the density that produces the maximum flow. The derivative at a particular point gives the space mean speed, which is the change in flow over the change in density in miles per hour [199, 12], shown in Equation (2.10):

$$u_s = qs \tag{2.10}$$

where  $s$  is the average space headway. The second relationship is between speed and flow. With flow on the x-axis and speed on the y-axis, we have the same flow-density curve, but rotated 90 degrees clockwise. As the flow increases from zero, the speed initially decreases very little, but gradually increases its rate of decrease as the flow continues to increase [199, 12]. Once the maximum flow is reached, the speed then starts to drop dramatically, as traffic begins to move into a congestion phase [199, 12]. As a result, it is desirable for transportation systems to operate at maximum flow, but in the real world, due to instability and congestion immediately beyond maximum flow, most hope to operate at just below maximum flow with room for error, at speeds slightly higher than those that occur at maximum flow. Since capacity - a rate - is the maximum output flow that can be achieved within a given time period, any attempts to put more vehicles into the system beyond capacity will result in congestion [199, 12].

## Flow Models

Given these fundamental parameters, we now elaborate on various traffic flow models, which explicitly describe the physical propagation of traffic flows within a pipe or across many pipes. These models can be broken down by their level of granularity in analysis, as there are (1) macroscopic, (2) mesoscopic, and (3) microscopic models [136]. Macroscopic models are based on fluid mechanics, and are the highest level of aggregation and thus the lowest level of detail [136]. Mesoscopic models are at a lower level of aggregation than macroscopic ones, and often use gas-kinetic theory in which driver behavior is considered [136]. Microscopic models are at the lowest level of aggregation and contain the most detail, describing minute interactions among vehicles in a pipe [136].

Macroscopic flow models are based on the continuity assumption in classic fluid mechanics, in which fluids should be treated as continuous media, such as a liquid in a pipe, rather than distinct particles [136]. Traffic is considered an inviscid but compressible fluid, and thus, flows, densities, and speeds are continuous variables, as noted by Maerivoet and Moor [136]. This leads to the canonical LWR model, which states that traffic as a fluid becomes less viscous as you move from congested to free-flow [223]. Such a notion implies a wave motion to traffic, as drivers that go from congested to free flow generate a fan of rarefaction waves [223]. The reverse of going from a low to high density regime results in a shockwave [136]. Daganzo builds on this model to distinguish different lanes and different types of drivers, which creates a capacity funnel [60]. This funnel shows that once a capacity drop occurs, the recovery to the capacity-flow regime cannot occur spontaneously, and thus it requires an exogenous mechanism [60]. Payne found that the LWR model is unsatisfactory in modeling real-time traffic flows because it did not account for stop-and-go waves, instead developing a higher-order model. He identifies three aspects in the momentum equation: (1) convection term, which describes how the space-mean speed changes due to the arrival and departure of vehicles at the time-space location, (2) relaxation term, which describes how

vehicles adapt their speed to the conditions dictated by the fundamental diagram, and (3) anticipation term, which describes how vehicles react to downstream traffic conditions [116]. However, Daganzo and the Berkeley school debunked these higher-order models in 1995, noting physical flaws in higher-order models and the lack of numerical stability in the solution space [136, 60].

Mesoscopic flow models look at aggregate characteristics like macroscopic ones, but they also consider such microscopic characteristics as driver behavior. One common mesoscopic model is the cluster model, which groups nearby vehicles together based on space-mean speeds rather than doing analysis on individual vehicles [136, 223]. Leonard and Ben-Akiva's worked on CONTRAM and DynaMIT, respectively, using this model [136]. Another common model is the headway distribution model, which is based on probability distributions of time headways and successive vehicles [136, 223]. Buckley and Branston developed semi-Poisson and generalized queuing models, respectively, under this umbrella. However, they assume homogeneous traffic flows, which is inadequate to describe real-world traffic dynamics [41]. Yet another common model is the gas-kinetic model, which is based on a phase-space density equation derived from Boltzmann's theory of gas dynamics [136, 60]. From the equation, we can get accelerations and interactions between the vehicles, as arriving and departing vehicles cause a change in the distribution of vehicle speeds. Despite the deeper granularity, mesoscopic models tend to fail at moderate to higher densities [136, 60, 12].

Microscopic models provide even greater granularity, giving explicit consideration to the interactions between individual vehicles within a traffic stream, employing characteristics like vehicle lengths, speeds, accelerations, and time and space headways. Car-following and follow-the-leader models, based on details of the interaction between two neighboring vehicles in a pipe, are the most widely known microscopic models [210]. The intelligent driver model (IDM) is then based on these theoretical frameworks, described later in this chapter. There are two stability criteria for car-following models: local and asymptotic [210]. Local stability describes how initial disturbances in the behavior of a leading vehicle affect a following vehicle. Asymptotic stability describes the behavior of the platoon of following vehicles, with instability occurring after the fifth following vehicle [210].

Microscopic models also take into account lane changing dynamics. Specifically, there are two types of lane changes: mandatory and discretionary [238]. Mandatory lane changes occur when the vehicle must take an exit off of the road, while discretionary lane changes occur for overtaking or voluntary movement [238]. In order to do a lane change, a check is made to see whether it is physically possible to merge into an adjacent lane, using gap-acceptance theory [238]. Within gap-acceptance theory, there are both optimal velocity models and psycho-physiological spacing models that dictate driver behavior [136]. The former manipulates the acceleration of a vehicle such that a vehicle's desired speed is selected on the basis of its space headway rather than just the speed of the leading vehicle. The latter creates spacing based on a driver perception threshold to speeds, speed differences, and space gaps.

Since microscopic models have greater granularity, its representation of real world dynamics is strong, but the computational complexity during simulation often becomes intractable.

As a result, traffic cellular automata (TCA) models have become more popular, as they advance time with discrete steps and in space that is coarse-grained. These models are flexible and efficient, but have low accuracy on a microscopic scale [136]. Brockfeld et al. benchmarked ten common microscopic traffic models for accuracy of predicted travel times, and found that the intelligent driver model and cell-transmission model performed best, with accuracies of nearly 83% [38].

Traffic modelers use three alternative types of traffic assignment models to predict the impact of travel demand on the network: (1) macroscopic, (2) mesoscopic, and (3) microscopic, in decreasing order of traveler aggregation and increasing order of granularity [145, 229]. Macroscopic models are based on the continuum assumption in classical fluid mechanics. The traffic flow is treated as continuous, similar to a flow of a liquid in a pipe, rather than that comprising of discrete vehicles [137]. These macroscopic models are useful in analyzing traffic systems covering a wide area, often across regions, and on highways where the overall speed dictates the macroscopic changes [123]. Unlike macroscopic models that assume a continuous vehicular flow on the road link (edge), mesoscopic models employ aggregated volume delay functions, by clustering a set of vehicles into packets and evaluating the movement of these clusters [123]. In contrast to these models, microscopic traffic simulation models provide even greater granularity, giving explicit consideration to the interactions between individual vehicles within a traffic stream and employing characteristics such as vehicle lengths, speeds, accelerations, time, and space headways [210].

## Simulation Tradeoffs

Regional-scale transportation modeling has been dominated by the macroscopic and mesoscopic models, due to their relative computational efficiency and familiarity [124]. However, one of the significant drawbacks of these simulators is their lack of granularity. They are limited by the accuracy of representing real-world vehicle dynamics, especially in congested regimes and for emergency scenarios [210, 16]. Traffic flow dynamics are naturally an outcome of the interaction of a many-particle system, where each particle exhibits different characteristics [236]. Only a microsimulation model can capture these intricacies of individual components and complex interactions with reasonable accuracy [133, 236, 210, 93]. However, microsimulation has a high computational cost due to the granularity required in simulating of vehicle movements. Hence, regional-scale microsimulation has generally been impractical [182, 123]. Although many traffic simulators exist, such as MATSim, SUMO, AIMSUN, Polaris, TRANSSIM, VISSIM, and DynaMIT, among others, these simulators are not designed to tackle large-scale traffic microsimulation efficiently [108, 125, 26, 14, 183, 165, 30]. As a result, techniques such as sampling a small fraction of the transportation demand are currently employed to achieve regional-scale traffic models in a reasonable amount of time and computational cost.

The next section describes a novel contribution to the microscopic simulator literature, MANTA, that carefully addresses the granularity vs. spatial extent vs. computational time tradeoff that exists.

## 2.2 Background

This section introduces a massively parallelized GPU implementation of a metropolitan-scale microsimulation engine - MANTA. MANTA is an agile regional-scale microsimulator capable of efficiently simulating over 7 million agents at a spatial scale as large as the San Francisco Bay Area, in under 10 minutes. First, we present the components of the simulation, then the mathematical theory and implementation of the simulator, followed by the results of a case study in the Bay Area. We then present the calibration and validation of the simulator, performance benchmarks, limitations and future work, and finally the conclusions.

Traffic modelers use three alternative types of traffic assignment models to predict the impact of travel demand on the network: (1) macroscopic, (2) mesoscopic, and (3) microscopic, in decreasing order of traveler aggregation and increasing order of granularity [145, 229]. Macroscopic models are based on the continuum assumption in classical fluid mechanics. The traffic flow is treated as continuous, similar to a flow of a liquid in a pipe, rather than that comprising of discrete vehicles [137]. These macroscopic models are useful in analyzing traffic systems covering a wide area, often across regions, and on highways where the overall speed dictates the macroscopic behavior [123]. Unlike macroscopic models that assume a continuous vehicular flow on a road link (edge), mesoscopic models employ aggregated volume delay functions, by clustering a set of vehicles into packets and evaluating the movement of these clusters [123]. In contrast to these models, microscopic traffic simulation models provide even greater granularity, giving explicit consideration to the interactions between individual vehicles within a traffic stream and employing characteristics such as vehicle lengths, speeds, accelerations, time, and space headways [210].

Metropolitan-scale transportation modeling has been dominated by the macroscopic and mesoscopic models, due to their relative computational efficiency and familiarity [124]. However, one of the significant drawbacks of these simulators is their lack of granularity. They are limited by the accuracy of representing real-world vehicle dynamics, especially in congested regimes and for emergency scenarios [210, 16]. Traffic flow dynamics are naturally an outcome of the interaction of a many-vehicle system, where each vehicle exhibits different characteristics [236]. Only a microsimulation model can capture these intricacies of individual components and complex interactions with reasonable accuracy [133, 236, 210, 93]. However, microsimulation has a high computational cost due to the granularity required in simulating the vehicle movements. Hence, metropolitan-scale microsimulation has generally been impractical [182, 123]. Although many traffic simulators exist, such as MATSim, SUMO, AIMSUN, Polaris, TRANSSIM, VISSIM, and DynaMIT, among others, these simulators are not designed to tackle metropolitan-scale traffic microsimulation efficiently [108, 125, 26, 14, 183, 165, 30]. As a result, techniques such as sampling a small fraction of the transportation demand are currently employed to achieve metropolitan-scale traffic simulation in a reasonable amount of time and computational cost.

This paper introduces a massively parallelized GPU implementation of a metropolitan-scale microsimulation engine - Microsimulation Analysis for Network Traffic Assignment (MANTA). MANTA is an agile metropolitan-scale microsimulator capable of efficiently sim-

ulating over 7 million agents at a spatial scale as large as the San Francisco (SF) Bay Area, in under 10 minutes. First, we present the components of the simulation, then the mathematical theory and implementation of the simulator, followed by the results of a case study in the Bay Area. We then present the calibration and validation of the simulator, performance benchmarks, limitations and future work, and finally the conclusions.

## 2.3 Components

The objective of this study is to perform a metropolitan-scale microsimulation of vehicular traffic of the SF Bay Area, incorporating individual trips on a typical workday morning. The microsimulator builds on the initial implementation by [87, 227]. In this section, the network generation, demand creation, routing, and simulation architectures are described in detail.

### Street Network

For the case study, we use the SF Bay Area, which includes nine counties. The street network is constructed from the OpenStreetMap (OSM) network within the polygonal hull of the counties in the metropolitan area using the OSMnx library [34]. The network contains all roads in the SF Bay Area, from large primary roads to tertiary streets. The OSM network includes intermediary points representing curves or bends in the road links, which are topologically unnecessary for network analysis [227, 34]. As these intermediary nodes do not represent a real intersection, they are removed from the network used for the microsimulation. The final network used for the microsimulation only includes nodes that represent physical intersections or dead-ends. The resulting network is a connected graph of the San Francisco Bay Area, where there exists a path from any node on the graph to any other node in the graph. There are no hanging nodes without a path. Figure 2.2 shows the full network of the SF Bay Area with 224,223 nodes and 549,008 edges. The number of lanes, length, and free-flow speeds for each edge are extracted from OSM data or imputed. The speed limit of each edge is taken from OSM, if available; if not, a free-flow speed limit is computed based on the number of lanes and the type of road. If the number of lanes is not available, then a recommended default value from OSM is used, depending on the type of road. For instance, a tertiary road without a specified number of lanes is given a default speed limit of 20 mph, and a motorway without a specified number of lanes is given a default speed limit of 57.5 mph.

### Demand

The origin-destination (OD) demand is derived from data generated by the Bay Area Metropolitan Transportation Commission (MTC) travel model [143]. For simplicity and considering the near-symmetric bimodality of morning and evening travel demand [143], this study only




Figure 2.2: The Bay Area Network with 224,223 nodes and 549,008 edges

considers the morning trips between 5 AM and 12 PM. The demand is further restricted to only automobile trips, which includes private automobiles, transportation network company (TNC) vehicles, and driving trips to transit stops. This demand does not consider public transit trips, such as buses, or large freight trucks, as the Bay Area MTC travel demand model is synthesized from only household travel data. While the data do not accurately reflect the exact real-world network congestion, leaving approximately 250K trips off the SF Bay Area road network for the 5 AM to 12 PM timeframe, the travel demand data remain comprehensive with approximately 3 million trips [143]

The OD pairs constructed from the MTC model are available at the granularity of traffic analysis zones (TAZ). TAZs are population density-based geographical areas typically larger than blocks but smaller than zip codes. The SF Bay Area consists of 1454 TAZs. The travel demand data is at the TAZ level (i.e., the OD pair represents a trip from an origin TAZ to a destination TAZ). The microsimulation requires an origin and destination node within its respective TAZ, which is achieved using a two-step process. First, all the nodes in the road network are mapped to their respective TAZ polygon. Second, for each TAZ-level OD pair, the origin-destination nodes are determined by sampling from a uniform distribution within the origin and destination TAZs, respectively. This process differs from [227], where origins and destinations are assigned to the centroid of their respective TAZs rather than being distributed across different nodes within the TAZ. The random assignment avoids unrealistic congestion at the centroids of the TAZs. The final OD demand has 3,269,864

travelers.

## Routing

After generating the network and the corresponding OD demand, the next step of the simulation is routing, where we compute the shortest-path between each origin and destination pair. Routing algorithms have been bottlenecks in many traffic models, requiring either significant pre-processing time or great computational cost [67]. The edge lengths, obtained from OpenStreetMap, represent the length of the road link in meters, which is used as the weight of the edge in the graph network for calculating the shortest path between nodes.

One of the significant contributions of this paper is the integration of a parallelized Dijkstra's priority queue single-source shortest-path (SSSP) algorithm, described in [247], in which only the OD pairs required in the simulation are computed. The priority-queue algorithm is parallelized with a hybrid MPI/OpenMP scheme, which allows for linear scaling with millions of agents on a cluster. An open addressing scheme-based hashmap is used to store key-value pairs of edge weights, which allows for faster update of the edge weights. This open addressing scheme improves the performance of hashmaps by 20%, providing quicker access to edges and connectivity. However, in this research, the road length is used as the edge weight and remains unchanged throughout the simulation. This feature will be useful in the future when the edge weights are updated dynamically based on traffic. A simulation with 3.2 million OD pairs routes on the large SF Bay Area network is calculated within 62 minutes on an Intel I9 processor with 2 threads per core and 14 cores per socket.

## Microsimulation

The microsimulation framework we adopt is an enhanced and extended version of the architecture developed by [87]. The vehicles move in discrete timesteps of  $\delta t = .5$  seconds, following the state of the art microsimulators today [72]. The simulation described in this paper models a typical morning workday from 5 AM to 12 PM. Each traveler in the OD demand is randomly assigned a departure time within this specified range by sampling from a normal distribution that roughly mimics the morning peak-hour behavior, with a peak around 7:30 AM and a standard deviation of 45 minutes. The departure time specified for individual vehicles in the simulation is presented in Figure 2.3.

At each timestep, the vehicle's travel time, position, and velocity are updated. MANTA employs a unique traffic atlas concept, akin to a texture atlas in the computer graphics community or a discretization step in signal processing. Each road segment of length  $l$  is discretized into  $\frac{l}{t_m}$  different compartments, where  $t_m$  is the length of each compartment, in meters, and is specified to be 1 meter in this implementation. Each 1 m compartment of the road is assigned to a specific byte in the computer memory. Hence, a road represented in the computer memory is a contiguous sequence of bytes. Each byte in memory can be occupied by at most one vehicle, and records the velocity of the vehicle and reflects their real position on the road lane [87]. Each road segment is represented as a new row in the 2D traffic atlas,






Figure 2.3: Departure times are chosen between 5 AM to 12 PM to model the morning hours. It follows a Gaussian distribution in which the bulk of the trips begin in the 8 AM to 9 AM hour.

and ultimately has a different number of bytes depending on its length. The 2D array layout of the road network requires less memory, making it parallelizable on a GPU rather than requiring a graph-based approach. The traffic atlas significantly reduces the computational cost of finding nearby vehicles, as it only involves looking up the status of neighboring cells in the memory array. The lookup scheme on a 2D grid to identify the vehicle’s speed and its neighbors’ speeds is parallelized with each thread on a GPU querying a specific block of memory address, which significantly speeds up the computation on the GPU. This approach varies from the traditional method, which requires checking the entire edge for neighboring vehicles.

In MANTA, the vehicular movement on an edge is dictated by conventional car following, lane changing, and gap acceptance algorithms [210]. The well-known Intelligent Driver Model (IDM), as shown in Equation (6.1), is used to control the vehicle dynamics through the network [213].

$$\dot{v} = a \left( 1 - \left( \frac{v}{v_o} \right)^\delta - \left( \frac{s_o + Tv + \frac{v\Delta v}{2\sqrt{ab}}}{s} \right)^2 \right) \quad (2.11)$$

where  $\dot{v}$  is the current acceleration of the vehicle,  $a$  is the acceleration potential of the vehicle,  $v$  is the current speed of the vehicle,  $v_o$  is the speed limit of the edge,  $\delta$  is the acceleration exponent,  $s$  is the gap between the vehicle and the leading vehicle,  $s_0$  is the minimum spacing allowed between vehicles when they are at a standstill,  $T$  is the desired time headway, and  $b$  is the braking deceleration of the vehicle [87, 227]. The exact position of each vehicle at a given timestep is computed by double integrating the acceleration value  $\dot{v}$ . Table 2.1 shows the range of possible values for the IDM parameters  $a$ ,  $b$ ,  $T$ , and  $s_0$ , which are derived based on the simulations from [213]. These parameters are calibrated for the microsimulation of the SF Bay Area using real-world data, with the process described in section 2.5.

Parameter	Value	Units
$a$	$\mathcal{N}(1, 10)$	$\frac{m}{s^2}$
$b$	$\mathcal{N}(1, 10)$	$\frac{m}{s^2}$
$T$	$\mathcal{N}(.1, 2)$	$s$
$s_0$	$\mathcal{N}(1, 5)$	$m$

Table 2.1: IDM parameter ranges, derived from [213]

MANTA is designed to be modular to incorporate different vehicle profiles and dynamics. In the future, these could include models for adaptive cruise control (ACC) and autonomous vehicles. Specifically, since the IDM leveraged by MANTA has been used in ACC systems in previous literature, it can be quickly adapted to accommodate for ACC in the future [119]. In ACC, the IDM maintains an appropriate relative distance to the lead vehicle, in contrast to standard cruise control systems, whose objective is to maintain only a target speed [119, 146]. Automated and platooned systems are also capable of cooperative adaptive cruise control (CACC), where there is communication among interacting vehicles [119]. Since the control system in each vehicle can retrieve information from the adjacent vehicles, such as the acceleration, velocity, and braking, these values can be stored as the IDM parameters in MANTA. The current setup uses the same values across all vehicles, based on calibration with real-world data (see section 2.5).

In addition to car following, vehicles can also change lanes within an edge. There are two types of lane changes: mandatory and discretionary [87]. Mandatory lane changes occur when the vehicle must take an exit off the road, while discretionary lane changes occur during overtaking or voluntary movements [87]. The lane changing model gives the vehicle an exponential probability from switching from a discretionary lane change to a mandatory lane change, as shown in Equation (6.2).

$$m_i = \begin{cases} e^{-(x_i-x_0)^2} & x_i \geq x_0 \\ 1 & x_i \leq x_0 \end{cases} \quad (2.12)$$

where  $m_i$  is the probability of a mandatory lane change for vehicle  $i$ ,  $x_i$  is the distance of vehicle  $i$  to an exit or intersection, and  $x_0$  is the distance of a critical location, which may be the position of a particular message sign (such as a final exit warning) to the exit or intersection (statically set to 1) [111, 236]. Intuitively, as the vehicle travels further along in a path, its probability of making a lane change to make a turn or exit increases.

Once a vehicle has decided to change lanes, the maneuver is performed if the lead and lag gaps are acceptable. A lane change dynamic involves interaction among three vehicles:

the merging vehicle  $i$ , the lead vehicle  $a$ , and the lag vehicle  $b$ . The critical lead or lag gap for a successful lane change is defined as the minimum distance to the following or lagging vehicle at which a lane change can be performed, respectively, as shown in Equation (2.13) and Equation (2.14).

$$g_{lead} = \max(g_a, g_a + \alpha_a v_i + \alpha_i(v_i - v_a)) + \epsilon_a \quad (2.13)$$

$$g_{lag} = \max(g_b, g_b + \alpha_i v_i + \alpha_b(v_i - v_b)) + \epsilon_b \quad (2.14)$$

where  $g_{lead}$  is the critical lead gap for a lane change,  $g_{lag}$  is the critical lag gap for a lane change,  $g_a$  is the desired lead gap for a lane change,  $g_b$  is the desired lag gap for a lane change,  $v_i$  is the speed of the merging vehicle  $i$ ,  $v_a$  is the speed of the lead vehicle,  $v_b$  is the speed of the lag vehicle,  $\alpha_i$  is the anticipation time of vehicle  $i$  attempting to change lanes in between vehicles  $a$  and  $b$  (in seconds), which have anticipation times  $\alpha_a$  and  $\alpha_b$ , respectively. Because drivers perceive distances and times differently, this anticipation time varies from vehicle to vehicle. The values of  $\alpha_i$ ,  $\alpha_a$ , and  $\alpha_b$  were chosen in the range  $[0.05, 0.40]$ , based on historically calibrated models [49]). If  $\alpha_i$ ,  $\alpha_a$ , or  $\alpha_b$  equals 0, this means the anticipation time of vehicle  $i$ ,  $a$ , or  $b$ , for vehicle  $i$  to make a lane change is 0 (i.e., the vehicle anticipates an extremely aggressive lane change with a low desired gap). Importantly, the anticipated gap is calculated based on the assumption that other drivers will maintain their current accelerations. For example, if the lag vehicle of the merging driver is decelerating, the anticipated gap will increase. Finally,  $\epsilon_a$  and  $\epsilon_b$  are the random components, each normally distributed with mean 0 and standard deviation 1, with units in meters. The values of  $g_a$  and  $g_b$  are a function of the speeds of the merging, leading, and lagging vehicles ( $i, a, b$ , respectively), but typically range between 1 to 5 meters at speeds below 25 mph and 5 to 10 meters at speeds above 25 mph [87].

The representation and modeling of intersections in this initial application of the traffic simulator is simplistic and not representative of diverse real-world dynamics at intersections. We consider two different types of traffic control. Case 1 traffic control is a flashing red light at each node, where only one vehicle can move into the intersection at a particular time. If the node contains  $n$  inbound edges and  $m$  outbound edges, the system will create a round-robin of the  $nm$  combinations for all cars to pass through the intersection based on their position in their lane queue [227]. Case 2 traffic control assigns every node as a green light, where all cars pass through the intersection with no delay. This is clearly not realistic for most nodes that have stop signs or traffic lights, but is plausible for nodes along highway interchanges. However, congestion on the edge itself is still modeled, as each edge has a finite capacity for vehicles and will not allow new vehicles from adjacent edges if the capacity is reached. The results of the simulations for these two cases are discussed in section 2.5.

## 2.4 OD and Routing Results

Preliminary travel patterns already emerge from the initial routing calculation. Figure 2.4 displays a log histogram of edge volumes across the network, showing that the bulk of edges




Figure 2.4: Histogram (log y-axis) showing the number of edges that see a particular vehicle count across the time range simulated. The Bay Bridge is the most heavily traversed link in the Bay Area, as it connects two major economic hubs: San Francisco and Oakland. The subset of edges representing the Bay Bridge and its necessary feeder edges (7 in total) sees a significant traffic volume of 30K trips, which has approximately 7.5K vehicles more than the next highest edges. Most edges see fewer than 100 vehicles in the timeframe.

have vehicle counts below 1000 and are traversed infrequently. Only a small fraction of the edges account for the majority of the most frequently traversed routes. Routes across the Bay Bridge are shown in Figure 2.5. Unsurprisingly, the Bay Bridge remains a unique outlier, as it accounts for a maximum volume of 31270 vehicles in the seven-hour duration. From [1] by AC Transit and ARUP, 41727 trips out of a total of approximately 4M trips traverse the Bay Bridge between 5 AM and 12 PM, representing 1% of all trips. This proportion of Bay Bridge traversals matches the proportion from the routing output at roughly 0.98% (31270 trips out of 3.2M total trips in the SF Bay Area).

## 2.5 Simulation Results

Infrastructure and scenario planning requires a high degree of accuracy in modeling the vehicle dynamics. This section highlights the calibration and validation techniques along with the microsimulator results. Previous studies have relied on vehicle counts, queue lengths at intersections, and vehicle speeds at loop detectors as ground truth data for calibration and validation [204].

In this work, we adopt a novel approach of calibration and validation using link-based speed data from the open-access Uber Movement project ([https://movement.uber.com/explore/san\\_francisco/speeds](https://movement.uber.com/explore/san_francisco/speeds)). Uber Movement uses GPS data from Uber vehicles to calculate speeds at particular snapshots in time. The edge speeds are calculated using the

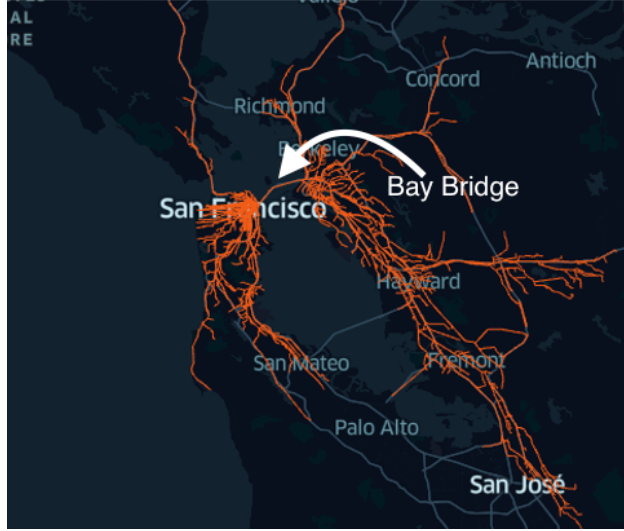


Figure 2.5: A network showing all possible routes through the Bay Bridge

GPS locations of vehicles and the map data of the street network. The Uber app records a vehicle’s location information every 1 to 2 seconds, including latitude, longitude, speed, course, and timestamp of the GPS location ping. Uber uses map-matching with the latitude and longitude of a particular GPS ping to determine when a driver enters and exits an edge. The duration a vehicle spends on an edge is calculated as the time difference between when the driver enters the street segment and when the driver leaves that segment. The edge speed is then calculated as the length of the segment divided by the time taken to traverse. Uber does not disclose the volume data, but notes that they ensure the validity of the data by setting a minimum threshold for observations.

## Calibration

Traffic microsimulators require calibration to real-world data to adequately represent observed dynamics across a wide range of network structures and conditions [26]. In the IDM, parameters  $a$ ,  $b$ ,  $T$ , and  $s_0$  are calibrated. The objective of the calibration process is to minimize the sum of the errors between every edge’s speed from MANTA and the Uber data (L1 norm). This optimization problem is specified in Equation (2.15),

$$\min \sum_{n=1}^N \left| \frac{\sum_{k=1}^K v_{k,n}}{K} - \bar{v}_{u,n} \right| \quad (2.15)$$

where  $v_{k,n}$  is the calculated velocity of vehicle  $k$  on edge  $n$ ,  $\bar{v}_{u,n}$  is the average Uber velocity of edge  $n$ ,  $K$  is the number of cars on edge  $n$ , and  $N$  is the number of edges that were successfully




Figure 2.6: The calibration process: average mean difference between Uber and MANTA speeds over time

matched between Uber’s street network and MANTA’s street network. Expanding further in Equation (2.16),

$$\min_{a,b,T,s_0} \sum_{n=1}^N \left| \frac{\sum_{k=1}^K [a(1 - (\frac{v_{k,n}}{v_{0,n}})^\delta - (\frac{s_0 + Tv + \frac{v\Delta v}{2\sqrt{ab}})}{s})^2]t}{K} - \bar{v}_{u,n} \right| \quad (2.16)$$

where  $t$  is the timestep (set as .5 seconds),  $a$  is the acceleration potential,  $b$  is the braking potential,  $T$  is time headway, and  $s_0$  is the linear jam distance.

Given the highly nonlinear nature of the objective function, a numerical method is used to optimize the IDM simulation parameters. We constrain the acceleration and deceleration potential,  $a$  and  $b$ , respectively, to  $[1, 10]$  meters per second squared, headway time  $T$  to  $[0.1, 2]$  seconds, and linear jam distance  $s_0$  to  $[1.0, 5.0]$  meters, and set the standard exponent of the IDM,  $\delta$ , to 4 [213]. A mini-batch gradient descent is then carried out across the entire simulation, with each iteration executing runs for 5 different sets of  $a$ ,  $b$ ,  $T$ , and  $s_0$ . We accumulated the sum of difference in speed between MANTA and Uber for all edges. The goal is to find the set  $\{a, b, T, s_0\}$  that minimizes this sum of differences. The set that produces the lowest mean difference is chosen as the nominal vector for the next iteration. Each parameter is then perturbed by a value chosen randomly, sampled from a uniform distribution, from  $[-1, 1]$  at the next iteration. Such a large range is used in order to produce meaningful differences across sets within the next iteration. The perturbation range then decreases by an order of magnitude at every iteration (e.g., iteration 3 uses  $[-.1, .1]$ , iteration 4 uses  $[-.01, .01]$ , etc.). The calibration process converges once the mean difference decreases below a desired threshold of .05 miles per hour, considering runtime limitations. As shown in Figure 2.6, the calibration process converges after five iterations.

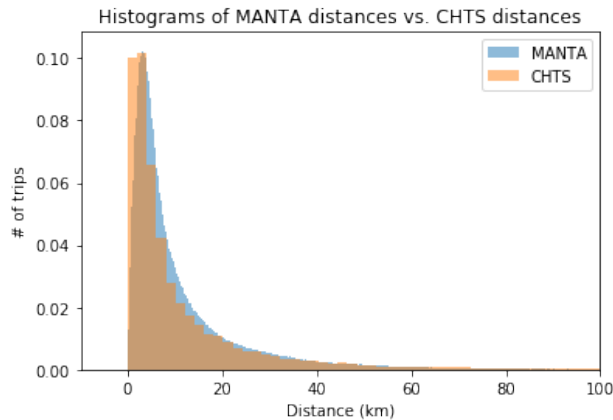


Figure 2.7: Comparison of trip lengths in MANTA versus California Household Travel Survey data. Median distance in MANTA is 6.46 km and in CHTS is 5.33 km.

Since the loss function in Equation (2.15) is non-convex across all of the calibrated parameters, this mini-batch gradient descent method will produce a local minimum and not necessarily a global minimum. When a successive order of magnitude decay, producing a larger perturbation range at every iteration, we observed that the gradient diverges, indicating a minima. Further research is required to improve the calibration method to approach global optimality.

## Validation

Validation is performed for both the routing algorithm and the traffic microsimulator. MANTA’s routing algorithm is validated by comparing the routes against the California Household Travel Survey (CHTS) data for the SF Bay Area [153]. Figure 2.7 presents the distances traveled by each vehicle for both MANTA and CHTS. The distribution of distances is heavily right-skewed, suggesting that most trips are fewer than 25 km. While CHTS data are sparse (69000 trips versus 3.2M trips in MANTA), we can still see similarities. MANTA estimates the mean distance traveled as 11.3 km, which is close to 13.5 km in CHTS. Median values are 6.46 km and 5.33 km in MANTA and CHTS, respectively. The 75th percentile distances are also similar, at 13.6 km and 13.7 km for MANTA and CHTS, respectively. The modest differences between the MANTA and CHTS routing data may be attributed to stochastic error from random sampling of O and D location within the respective TAZs in MANTA simulation.

The validation of the traffic microsimulator involves comparing the MANTA outputs to Uber Movement distributions at specific timeslices. In particular, using Q2 Uber Movement data from 2019, we compare results of the MANTA simulations and the Uber data for 95,510 edges, or 17%, of the total edges in the SF Bay Area network.

We enhanced the IDM to include varying maximum speed limits for each individual to better reflect real-world vehicular behavior. In the IDM,  $v_0$  represents the free-flow velocity of a vehicle on an edge, typically the speed limit of each edge from OSM or from a standard convention. However, in order to mimic the variance of driving patterns across travelers, each driver's maximum possible speed per edge,  $v_0$ , is sampled from a Gaussian distribution centered around the edge's predetermined speed limit with a standard deviation of  $2\sigma_s$ , where  $\sigma_s$  is the standard deviation of vehicle speeds, obtained from the Uber data, at each speed limit  $s$ . Every vehicle thus has a slightly different maximum allowable speed on each edge it traverses.

We compare the distribution of speed on different edges between MANTA simulation and Uber data. For the simulation run between 5 AM - 12 PM, we investigate the difference in behavior at two different time periods: 5 AM - 6 AM, a less congested time period, and 8 AM - 9 AM, a more congested time period. Within each time period, we look at the speed distribution curves at different speed limits. For instance, Figure 2.8 and Figure 2.9 show the speed distributions from MANTA on edges with 35 mph compared to the Uber Movement data on those same edges, at the representative time periods. As expected, both MANTA and Uber average speeds are higher between 5 AM and 6 AM (less congested) than those between 8 AM and 9 AM (more congested time period).

Figure 2.11 and Figure 2.12 show the average speeds of MANTA and Uber across all speed limits between 5 AM and 6 AM. At low-speed edges ( $< 30$  mph), MANTA simulation speeds are approximately 5 mph slower than Uber's real-world data. This indicates that the congestion effects are larger at lower speeds in MANTA. The Uber speeds also reflect that, in the real-world, many drivers tend to go above the speed limits more so on edges with lower speed limits than they do on edges with higher speed limits. For edges with speed limits above 30 mph, MANTA estimates may be higher or lower than the Uber estimates. This suggests that improvements can be made in both calibration and in modeling the individual behavior of drivers with respect to speed limits.

Figure 2.13 shows the distribution of speeds for the 8 AM - 9 AM timeframe. Unlike the less congested 5 AM to 6 AM timeframe, MANTA simulation speeds are equal to or slower than Uber's real-world data across all speed limits. This indicates that the IDM in MANTA may be overly sensitive to congestion effects.

Comparing the 5 AM - 6 AM timeslice with 8 AM - 9 AM in MANTA, the average speeds estimated in the early morning time period in general are higher by 3 to 9 mph across all speed limits, with the greater differences being on edges with higher speed limits. This intuitively suggests that roads with higher speed limits, such as highways, see less traffic at the early morning hours, and thus vehicles can travel at higher speeds due to the lack of congestion and lack of stoppage. However, roads with lower speed limits do not allow for much higher speeds regardless of the time of the day. This is likely due to the presence of frequent intersections in the city. The Uber data across the two timeslices also reflect this difference.






Figure 2.8: 5 AM - 6 AM




Figure 2.9: 8 AM - 9 AM

Figure 2.10: Kernel density plot comparing the MANTA and Uber distributions at 35 mph




Figure 2.11: Average MANTA speeds across all speed limits [5 AM - 6 AM]. The means and standard deviations are shown in parentheses.

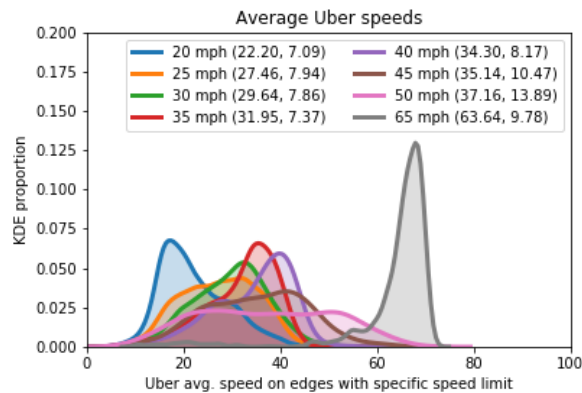


Figure 2.12: Average Uber speeds across all speed limits [5 AM - 6 AM]. The means and standard deviations are shown in parentheses.

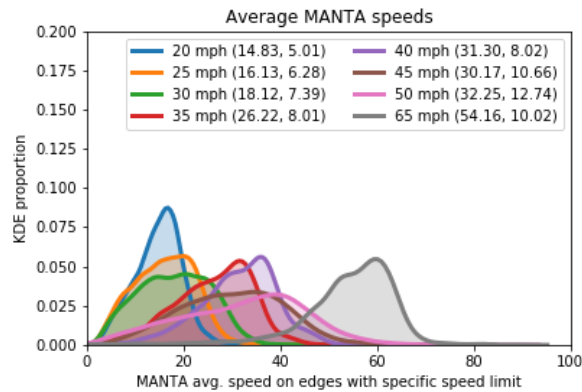


Figure 2.13: Average MANTA speeds across all speed limits [8 AM - 9 AM]. The means and standard deviations are shown in parentheses.




Figure 2.14: Average Uber speeds across all speed limits [8 AM - 9 AM]. The means and standard deviations are shown in parentheses.




Figure 2.15: Average MANTA speeds across all speed limits [8 AM - 9 AM] in the red light case. The means and standard deviations are shown in parentheses.

## Red light / green light cases

In this study, we adopt a basic intersection model and consider two different conditions: where every node is either a flashing red light or a green light. In the flashing red light scenario, every vehicle is designed to stop at the intersection for 2 seconds before proceeding, similar to a stop sign, which can result in vehicles backing up and subsequent queue spillback. In the green light scenario, every vehicle can immediately access the intersection and proceed with its next move. Figure 2.15 shows the distribution of average speed across different speed limits. Between 5 AM - 6 AM, the average speed is 17.5 mph, while the speed decreases to 12.9 mph in the 8 AM - 9 AM timeslice. The reduction in speed between 8 AM and 9 AM suggests increased congestion, in comparison to the free-flowing traffic in the early morning between 5 AM - 6 AM.

When every node is a green light, the average speed across all speed limits for the 5 AM - 6 AM time period is 24.5 mph. The average speed across all speed limits decreases to 17.8 mph for the 8 AM - 9 AM time period. Each speed limit's average speed is shown in Figure 2.13. The average difference in speed across all speed limits between the early morning timeslice (Figure 2.11) and the 8 AM - 9 AM peak hour timeslice (Figure 2.13) in the green light case is 4.6 mph. For the red light condition (whose figures are not shown out of brevity), it is 6.7 mph. The deltas between the two timeslices, as well as the absolute speeds, highlight notable differences in the traffic behavior between the two timeslices. Specifically, the average speeds in both timeslices under the red light condition is about 5 mph lower than the green light condition. Such low speeds are unsurprising given that every vehicle must stop and wait its turn in the intersection queue. Since the IDM parameters have been tuned to the real-world Uber data, which is better represented by the green light scenario, the red light scenario does not match the Uber data as closely as the green light scenario does.




Figure 2.16: Fit to lognormal distribution for 20 mph speed limit in green light scenario (case 2)




Figure 2.17: Fit to normal distribution for 45 mph speed limit in green light scenario (case 2)

Notably, in Figure 2.13, the lower speed limits' distributions tend to be right-skewed, following a lognormal pattern, while the distributions at higher speed limits become more centered and follow a normal distribution. Snapshots of these phenomena are shown in Figure 2.16 and Figure 2.17.

## 2.6 Performance benchmarks

This section describes the computational performance of the two core components of MANTA: routing and the microsimulator engine.

### Routing performance

In our network of approximately 225K nodes, 550K edges, and 3.2M OD pairs, the SSSP routing algorithm carries out the computation of all OD pairs' routes in approximately 62 minutes on a single node. Figure 2.18 shows the time-required to run up to 1 million agents on a distributed compute cluster utilizing both MPI and OpenMP parallelization schemes. Figure 2.19 shows that the strong scaling results of the routing algorithm matches the theoretical scaling up to 1024 cores for routing 1 million agents. In comparison to existing routing algorithms, such as the heuristic-based Ligra [190] and iGraph [57], the priority-queue based Dijkstra is 2.2% and 55% faster, respectively, on a single node. The priority-queue Dijkstra algorithm also has higher effective CPU usage of 94.1% with an average RAM usage of 4.81 GB.

### Microsimulator performance

The computational performance of the MANTA simulator is compared with Simulation for Urban Mobility (SUMO) and JDEQSIM, a parallelized alternative available in MATSim, two well-known open-source simulators in transportation. The simulation of the SF Bay Area network and the demand between 5 AM - 12 PM are used for the comparison exercise. SUMO offers two options to build the network: one that contains internal links or lanes within intersections, and one that does not contain internal links [125]. Considering MANTA's simplified intersection model, the SUMO model without internal links is the most appropriate comparison. The SUMO model with internal links is also included for completeness.

Section 2.6 shows the runtime comparison of MANTA against SUMO and JDEQSIM. The table also indicates when the results are linearly extrapolated, due to the inability to complete simulations in a reasonable time. Extrapolating the simulation runtime linearly, MANTA performs nearly 27000x faster than SUMO. MANTA carried out the full microscopic simulation of 3.2M trips at .5 s timesteps in 4.6 minutes, while SUMO's simulator is estimated to take nearly 87 days, linearly extrapolated from the initial run of 194 minutes for 5000 trips. SUMO also has a mesoscopic simulator, which requires approximately 29 hours (1740 minutes) for the SF Bay Area simulation.

A primary reason for such a dramatic difference in runtimes is that typically SUMO uses a traffic assignment model for routing. When the routes are fixed, as in this example, SUMO sees undesired jamming, as many roads are not filled to their capacities while other roads are filled excessively. The resulting congestion increases the simulation time in SUMO to achieve equilibrium. Unlike SUMO, MANTA is a dynamic model and does not perform equilibrium traffic assignment. In other words, MANTA does not minimize the total travel




Figure 2.18: Time required to route agents using priority-queue Dijkstra algorithm for the SF Bay Area network on distributed computing environment (MPI + OpenMP) parallelization. Tests were run on 32 nodes with Intel Xeon Skylake 6142 processors.

time of the entire system, but instead assumes that each driver will take the shortest route based on distance. Notably, SUMO’s microsimulation does not support parallelization; only the routing algorithm is parallelized, which is not germane for this comparison.

JDEQSIM is a discrete event-based mesoscopic simulator that uses event handling to communicate every person’s activity to the rest of the Behavior, Energy, and Autonomy Modeling (BEAM) platform[229]. The event handler manages billions of activities and events (specifically when vehicles enter and exit edges), which produces a significant overhead in the generation and synchronization of events across the threads. MANTA, on the other hand, is a time-based simulator that does not have overhead from the constant generation of events. In addition, the texture mapping of MANTA is optimized for fast GPU array manipulation, which yields significant speedup compared to the CPU implementation in JDEQSIM [229].

Figure 2.20 shows the comparison of runtimes between JDEQSIM and MANTA. The JD-




Figure 2.19: Speedup of priority-queue routing algorithm for the Bay-Area network on distributed computing environment (MPI + OpenMP) parallelization. Tests were run on 32 nodes with Intel Xeon Skylake 6142 processors.

EQSIM runtime is approximately 6.6 minutes, on average over 50 runs, and is comparable to MANTA's runtime of 4.6 minutes. The GPU parallelized traffic microsimulation in MANTA is 43% faster than aggregated simulators such as JDEQSIM. In comparison to the SUMO microsimulation, MANTA is several orders of magnitude faster. Considering the finer level of behavioral granularity achieved by MANTA at the runtime of the mesoscopic JDEQSIM, these results clearly demonstrate the applicability of MANTA for metropolitan-scale traffic microsimulations.

Other parallel microsimulators exist as well, including [46, 27, 151], but they either require expensive supercomputing facilities or carry out simulations on smaller networks with longer computation times.



Simulator	Time (mins)	Type
MANTA	4.6	Full
SUMO meso simplified (MeS)	1620	Full
SUMO micro simplified (MiS)	114858	Lin. extrap.
SUMO meso advanced (MeA)	1740	Full
SUMO micro advanced (MiA)	123500	Lin. extrap.
JDEQSIM	6.6	Full

Table 2.2: MANTA’s runtimes compared to SUMO and JDEQSIM. Full implies that the entire simulation was able to complete. Lin. extrap. implies that only part of the simulation was able to complete and the full time was linearly extrapolated from this preliminary time.




Figure 2.20: Simulator runtimes (log scale y-axis) across different simulators. MANTA performs slightly better than the parallelized mesoscopic JDEQSIM and is on the same order of magnitude. MANTA performs significantly better than the mesoscopic version of SUMO with either the simplified (MeS) or advanced intersection modeling (MeA). The microscopic version of SUMO with simplified intersections (MiS) and advanced intersections (MiA) could not be run completely, and thus times were linearly extrapolated, reflecting that it would take tens of days to complete.

## 2.7 Limitations

The traffic microsimulation in MANTA achieves significant advances in computational performance using metropolitan-scale networks and demand, but important limitations remain. The first limitation is the use of simplified intersection modeling. A more accurate intersection modeling will produce precise travel times and a better representation of the vehicle dynamics.

The second limitation is the demand profile. This work uses a synthetic Bay Area MTC 2017 travel model that represents the daily demand in five large time blocks and carries out a static traffic assignment. A more realistic model could integrate a dynamic travel demand model, such as ActivitySim, with MANTA.

The modular structure of MANTA offers the ability to vary different components of the network analysis, such as routing and vehicular dynamics. MANTA currently can accommodate different routing algorithms, such as Dijkstra, A\*, and Contraction Hierarchy. In addition, while MANTA currently uses the Intelligent Driver Model, it has the functionality to leverage other driver models. Incorporating dynamic routing, where the edge weights are based on travel times on the edge rather than the length of the edge, will improve the predictive accuracy of near-real-time simulations, such as evacuations.

## 2.8 Conclusions

This paper presents a novel traffic microsimulator, MANTA, that addresses the challenges of accurate traffic microsimulation at the metropolitan-scale. MANTA is highly efficient and is capable of simulating real-world traffic demand with a fine level of granularity on very large-scale networks. The runtime efficiency of MANTA is achieved by efficiently coupling a distributed CPU-parallelized routing algorithm and a massively parallelized GPU simulation that utilizes a novel traffic atlas to map the spatial distribution of vehicles as contiguous bytes in memory. The capability of MANTA is demonstrated by simulating a typical morning workday of the nine-county SF Bay Area network with 550K edges and 225K nodes, and approximately 3.2M OD pairs. The routing calculations are completed in 62 minutes, and a simulation of 7 hours from 5 AM to 12 PM with .5 second timesteps is completed in 4.6 minutes. This is several orders of magnitude faster than the state of the art microsimulators with similar hardware. Achieving compelling performance in both efficiency and accuracy, MANTA offers significant potential for fast scenario planning in both short- and long-term applications in metropolitan and metropolitan-scale analysis.

## Chapter 3

# Machine Learning and Emergent Mobility: Advances in Transportation Simulation and Modeling

This chapter presents emerging contributions to simulation and modeling, with a specific focus on the use of machine learning (ML) techniques to augment existing frameworks. The chapter concludes by describing how emergent mobility modes and new technologies are altering how we approach transportation simulation and modeling more broadly.

### 3.1 Machine Learning in Simulation

Machine learning (ML) and deep learning techniques have gained increased traction over the past decade, specifically in technical problems requiring what has been deemed “artificial intelligence.” Leveraging large-scale data for prediction and classification, machine learning has pervaded disciplines of inquiry long thought unrelated to computer science, including city planning and transportation [232]. Advanced ML applications have been deployed both in simulation and in real-world remote systems to understand and control mobility for better urban outcomes. Indeed, deep learning and computer vision using convolutional neural networks has been leveraged in sense-and-avoid systems deployed in emergent mobility technologies, such as AV and UAM [4, 69, 83, 105, 66]. Deep reinforcement learning algorithms, including policy gradients and Q learning, have been used in remote traffic control and optimization [232, 13, 69, 234, 86]. Traffic flow prediction has also benefited from ML, as nonparametric algorithms and neural networks have been utilized to perform traffic state estimation as well as forecasting in a variety of geographies [202, 207, 149, 115].

Given the critical nature of understanding the impact of such emergent mobility technologies as AV and UAM on urban outcomes, which include sustainability, equity, and economic productivity, leveraging ML in unique and innovative ways will prove useful for modeling and simulating more accurately travel behavior dynamics. The first section provides a brief

overview of machine learning and its theoretical foundation. The second section highlights the literature in traffic flow prediction that leverages machine learning methods. The final section presents open areas of research in urban simulation that may benefit from different ML models.

## Overview of ML

The concept of machine learning is derived from a data-driven approach to understanding systems and phenomena. In traditional programming, the CPU often takes as inputs data and the program itself, producing an output of the program performing particular functions on the data [63]. However, machine learning differs in both input and output. Specifically, the CPU now takes data *and* the output as inputs, instead producing the program as its output [63]. In other words, the CPU is determining the program or function that best represents the data based on output feedback [63, 113].

As a result, every machine learning algorithm has three major components: (1) representation, (2) evaluation, and (3) optimization. Representation is the broad term for the model generated and used [63, 113]. Each model varies from another in the way it learns from the data. Evaluation is necessary to determine the effectiveness of the model in carrying out the objective task. Optimization is important in improving the model once it has been initially evaluated.

It is also critical to note the various types of learning, which use different algorithms for particular applications. The first type of learning is supervised learning, which requires that ground truth data exist to confirm that the algorithm is learning from known data [63, 174]. In other words, there is a correct answer for the data during training. The second type of learning, and the converse, is unsupervised learning, for which the data do not have any labels on the input data [63, 113]. This inherently makes unsupervised learning more exploratory in nature, and its primary goal is to find structure or patterns in the data that perhaps may be obvious to human intuition. However, unsupervised learning is not as prevalent due to challenges in application [113]. The final type is reinforcement learning, in which an agent carries out a task in an unknown environment based on rewards and punishments [128]. Reinforcement learning is a key contributor to the artificial intelligence generally seen in robots and humanoids.

## Optimization

Developing a useful model is perhaps the most critical component of machine learning. In order to generate a strong model, the mathematical theory behind machine learning relies on optimization theory. Specifically, in all types of learning, the model or learner must learn from input data in order to produce a particular output, whether in a prediction, classification, or an agent-based task such as moving toward or avoiding an object. However, it must learn based on optimizing for a particular objective, such as minimizing the error between the input and output. For example, in a supervised learning context, when the learner develops

a prediction  $Y_i$ , it is optimal that this prediction is as close to a known ground truth output  $Y_i'$ , and thus the goal of the model is to minimize the error by calculating the appropriate parameters,  $\alpha$ ,  $\beta$ , &  $\delta$ , of the model that will do so. The error minimization becomes the cost or loss function, which is convex, as shown in Equation (3.1):

$$J(\alpha, \beta, \delta) = \frac{1}{N} \sum_{i=1}^N |\alpha + \beta X_i + \delta X_i^2 - Y_i| \tag{3.1}$$

and thus the objective is:

$$\min_{\alpha, \beta, \delta} J(\alpha, \beta, \delta) \tag{3.2}$$

In order to minimize  $J$ , first-order optimization must be carried out by calculating the partial derivatives of  $\alpha$ ,  $\beta$ , &  $\delta$  with respect to  $J$ . From Newtonian calculus, the next step is to set each parameter's partial derivative to 0 in order to find its optimal value, with an example in Equation (3.3) [113]:

$$\frac{\partial J}{\partial \alpha} = 0 \tag{3.3}$$

While this procedure is tractable even by hand for low dimensional problems, in real-world applications, the cost function can have thousands or even millions of parameters that must be optimized to our objective. As a result, it becomes computationally intractable to solve directly for each of these parameters. To overcome this challenge, gradient descent, an iterative algorithm designed to search for the optimal values for each parameter, is frequently used in machine learning and deep learning [31].

In gradient descent,  $\alpha$ ,  $\beta$ , &  $\delta$  are initialized randomly, and at each iteration, the partial derivatives (or gradients) are calculated simultaneously such that their values are updated, eventually leading to a global minimum if the function is convex, as shown in Equation (3.4) and ????:

$$\alpha < -\alpha - R \frac{\partial J}{\partial \alpha} \tag{3.4}$$

$$\beta < -\beta - R \frac{\partial J}{\partial \beta} \tag{3.5}$$

$$\delta < -\delta - R \frac{\partial J}{\partial \delta} \tag{3.6}$$

where  $R$  is the learning rate, an adjustable hyperparameter that can determine the speed and accuracy of the convergence. In particular, a small  $R$  may converge slowly but a large  $R$  may overshoot the minimum or may fail to converge altogether. A particular version of gradient descent known as the Adam optimizer is often used in deep learning in order to find the optimal weights and biases for each layer. Adam optimization often performs better

than standard gradient descent due to the existence of saddle points, which are not optima but still have gradients equal to 0; as a result, the standard gradient descent algorithm may stay stuck at that value despite its not being a minimum of any sort.

## Loss Functions

While in nonparametric models the loss function is often an error minimization such as root-mean-squared error, in many neural network architectures used for classification, the training will be driven by errors in distributions [128, 63]. Shown in Equation (3.7), the cross-entropy function minimizes the distance between two probability distributions.

$$E = - \sum Y_i^* \log(Y_i) \tag{3.7}$$

where  $Y_i^*$  represents the actual probabilities while  $Y_i$  is the computed probabilities.

Cross-entropy loss increases as the predicted probability diverges from the correct classification, and thus, a perfect model would have a log loss of 0. Indeed, minimizing cross-entropy loss is equivalent to performing maximum likelihood estimation, assuming classification labels are independent, as shown below [184]:

$$\begin{aligned} \hat{Y}_i &=_{\hat{Y}_i} \prod_{i=1}^N P(\hat{Y}_i = Y^*i) \\ &=_{\hat{Y}_i} \sum_{i=1}^N \log P(\hat{Y}_i = Y^*i) \\ &=_{\hat{Y}_i} \sum_{i=1}^N \sum_{j=1}^M P(Y_i^* = j) \log P(\hat{Y}_i = Y^*i) \\ &=_{\hat{Y}_i} \sum_{i=1}^N - H(P(Y_i^*), P(\hat{Y}_i)) \\ &=_{\hat{Y}_i} \sum_{i=1}^N H(P(Y_i^*), P(\hat{Y}_i)) \end{aligned} \tag{3.8}$$

Models in nonparametric machine learning, deep learning, and deep reinforcement learning utilize different optimization techniques, and specifically gradient descent, for different applications. However, virtually all contexts require an objective function, such as cost or likelihood, that must be minimized or maximized, respectively.

## Traffic Flow Prediction

ML frameworks have recently been used extensively in traffic flow prediction, which has become an urgent issue in transportation policy and planning. Modeling vehicle movements

over time is critical in understanding travel behavior, infrastructure provision and traffic control, sustainability impacts, and providing real-time routing.

### Nonparametric Techniques

Traffic flow prediction specifically attempts to estimate the number of vehicles given a specific spatial region and time interval, such as a road segment in a peak commute hour. Prediction depends on historical and real-time data, which is often collected from GPS data from smartphones and its associated applications, loop detectors within roads, radars, and street cameras [135, 231]. Previous traffic models have been shallow in using empirical data, and thus with greater computing capabilities, leveraging the plethora of traffic data may make it more accurate. Huang et al. noted that many existing techniques to predict traffic flow have used time-series [110]. This parametric time-series approach utilizes variations of the autoregressive-moving average, which focuses on the temporal variation of traffic flow to predict future traffic flows [135, 110, 168]. However, it assumes linear relationships between major macroscopic parameters such as density, flow, and speed with their future values, which inherently limits its representativeness in the real-world, particularly considering nonlinear incidental traffic phenomena such as accidents, construction, and shockwaves [110, 59].

As a result, nonparametric methods such as k-nearest neighbors (kNN) have become popular due to their ability to learn any functional form of the data, rather than being limited to parametric relationships such as linearity [200]. kNN is a supervised learning algorithm in which labeled training data  $(x, y)$  are given, and the goal is to find the relationship between  $x$  and  $y$  by learning a function  $h$ , such that  $h : X \rightarrow Y$ . By finding  $h$ , given an unknown  $x$ , we can predict  $h(x) = y$  for that value of  $x$  [113]. Since the objective in kNN is to find data points that resemble the data  $x$  as closely as possible, a distance, or similarity metric  $d$ , must be calculated from  $x$  to each point. Many distances can be used, but one of the most common is the Euclidean distance, as shown in Equation (3.9):

$$d(x, x') = \sqrt{(x_1 - x'_1)^2 + (x_2 - x'_2)^2 + \dots + (x_n - x'_n)^2} \quad (3.9)$$

Then, if each point's  $d < \epsilon$ , where  $\epsilon$  is a cross-validated threshold, then it will be part of a set  $A$  [113]. The conditional probability is then calculated for a class, or the fraction of points in  $A$ , which, in the case of traffic flow prediction, may be congested or free-flow, for instance. The conditional probability is calculated in Equation (3.10) [63]:

$$P(y = j|X = x) = \frac{1}{K} \sum_{i \in A} I(y^{(i)} = j) \quad (3.10)$$

The highest conditional probability is then chosen as the classification for the input  $x$ . Several works use a kNN approach to classify road segments based on other road segments whose characteristics include volume, speed, and occupancy rate [245, 163].

Other nonparametric machine learning techniques that have been used include support vector machines and nonlinear regression [110, 114, 200]. It is also crucial to note that while

model accuracy is important, it is also equally important to understand the environment, whether in space, time, or even required resources for a particular geography, in which the model is developed and utilized [195].

## Neural Networks

In addition to nonparametric machine learning approaches, deep learning has become increasingly more expedient in traffic flow prediction due to the ability of neural networks to understand patterns in nonlinear inputs and also due to the availability of large-scale datasets [135]. Much research has used multilayer neural networks to predict traffic flows based on such data as GPS trajectories and static loop detectors [135, 168, 114].

The building block neural network is the feedforward neural network, also known as a multilayer perceptron. A feedforward neural network architecture contains several steps. Typically, the methodology requires several neural layers, the first of which is the input layer of neurons and the following layers are hidden layers of neurons [128]. Each neuron in a neural network carries out a weighted sum of all its inputs and then adds a bias; the value of the output determines whether the neuron should ‘fire’ or ‘activate’ [128, 184]. Since the goal in both classification and regression problems is to determine a function of the inputs,  $f(x_1, x_2, \dots, x_m)$  that can accurately explain the observed outputs  $(y_1, y_2, \dots, y_m)$  [128], the function,  $f$ , is the nonlinear activation function and is the decision rule for activating the neuron. Many activation functions exist, such as the sigmoid, softmax, rectified linear unit (ReLU), and hyperbolic tan [63]. In the ReLU, shown in Equation (3.15), the activation function is thresholded at zero. Compared to the sigmoid, the ReLU does not have computationally intensive operations like exponentials. In addition, the sigmoid also suffers from vanishing gradients since it squashes all values between 0 and 1 [184].

$$f(x) = \max(0, x) \tag{3.11}$$

Despite these benefits of the ReLU, the decision of activation functions depends on the context, as binary classifiers may use simple step functions while less constrained functions, such as the ReLU or softmax, can provide a range of activations.

Other popular neural network architectures include the recurrent neural network (RNN) and convolutional neural network (CNN). RNNs take as inputs not only the current input, but also what they have “perceived” previously in time, essentially using the output at time  $t - 1$  as an input to time  $t$ , along with the new input at time  $t$ ; this is known as back-propagation. However, simple RNNs that must find connections between the final output and data several timesteps before are limited, since there is often an exponential number of multiplications that occur within the hidden layers of the network. Similar to one of the sigmoid function’s drawbacks, this creates derivatives that will vanish, which makes it difficult for processors to compute and for networks to learn [128]. For this reason, long short-term memory (LSTM) architectures, a subset of RNNs that maintain gradients through the back-propagation process, are often used, since they preserve the errors in a gated cell. RNNs



may be useful in analyzing trajectory data in a specific street network over long timespans, perhaps to develop long-term congestion models.

In CNNs, one neuron does a weighted sum of the input vector across only a subset of the input vector, sequentially choosing different shifted subsets of the input vector, similar to convolution in signal processing [160]. Each neuron reuses the same weights, while in feedforward networks, each neuron has its own set of weights [168, 128]. Typically, a number of layers are used, primarily because the initial layer represents the high level features of the input vector while subsequent layers represent more detailed features, reflected by more filters in those layers [184].

Finally, each neuron will be a weighted sum of all of its inputs with the bias, which will then become the input to a nonlinear activation function. The activation function in this last layer, particularly in CNNs, is typically the softmax, as shown in Equation (3.12):

$$\text{softmax}(L_n) = \frac{e^{L_n}}{\|\mathbf{e}^L\|} \quad (3.12)$$

where  $L_n$  is the weighted sum of all inputs and biases, and the softmax itself is the output of each neuron. In a standard classification setting, the output layer of the neural network must have  $n$  output neurons, where  $n$  is the number of possible classification types. For instance, in the example of classifying the road segment as either congested or free-flow, the output layer will have 2 output neurons.

### Overfitting

While there are substantial differences in each ML model’s methodology, in almost all supervised machine learning techniques, overfitting on known training data is often a significant issue. Specifically, since the learned model is based purely on training data, the model’s accuracy on the training data can be close to 100%, but the model may be unable to produce accurate results when presented with unknown test input data. In nonparametric models, cross-validation, randomization, and regularization of data are often ways to minimize noise and enhance generalizability [63]. In a sample loss equation in Equation (3.13), regularization penalizes a greater number of parameters, hoping to simplify the model away from the exact training data.

$$J = (y - \hat{y})^2 + \lambda w^2 \quad (3.13)$$

where  $J$  is the cost function,  $y$  is the actual data,  $\hat{y}$  is the predicted data, and  $\lambda$  is the regularization parameter that is essentially determined by the degree of dropout. Deep learning also uses regularization, manifested as a dropout layer in neural networks. Since overfitting is caused by the model being excessively tuned to only the training data and not being broadly generalizable to new test data, the dropout routine drops a random set of neurons by setting them to zero. This will cause robustness in the system in such a way that the classification or regression can still be made accurately without these neurons, which would be more specific to the actual input data anyhow. Another approach in deep

learning to battle overfitting is applying a maxpool layer, which downsamples the input. Since the input layer produces a value that shows that there is a specific feature that exists in the input, maxpooling allows that feature's exact location not to be as important as its location relative to other features. This not only can decrease overfitting but also reduce the computational burden.

## 3.2 An Example: Traffic Control Inference using Deep Learning

### Motivation

Traffic flow prediction and simulation are broadly important for transportation policy and planning. However, one of the glaring issues with existing traffic simulators is a lack of understanding of traffic control in cities. Several cities possess manual records of their many intersections, but they are often neither accessible nor comprehensive. Knowing all traffic control in a metropolitan area has been a challenging point for planners and travel modelers. Only certain cities have organized databases that outline the type of traffic control at each intersection, but even those, such as the one in San Francisco, is relatively sparse. In addition, since the traffic control that does exist is within only one city, it is not sufficiently comprehensive to cover many origin-destination pairs of travel demand, which are often across cities within a metropolitan area. Google Maps has also been incorporating traffic control into their maps and traffic forecasts, but their approach and the geographies covered remain unknown.

For this reason, instead of manually parsing inexact and inadequate databases, it is proposed to *infer* traffic control using more easily obtainable data and with advanced deep learning. Image recognition and computer vision techniques using convolutional neural networks (CNNs), highlighted by Zhang, Olah, Hofleitner, and Krizhevsky, serve as inspirational models [246, 157, 158, 106, 126]. Gao et al. and Liang et al. have used machine learning and deep learning techniques to optimize traffic signals, and Lv et al. are one of only a few groups who have used deep learning to understand traffic flow [135, 86, 131]; however, there has been little work that has used deep learning to infer traffic control for simulation. In this section, advanced deep learning techniques are used on granular travel demand data to infer traffic control, an aforementioned limitation of MANTA's current implementation. Los Angeles is used as a case study, with the greater goal of extrapolating the method to other cities for which we have mobility trajectory data.

### Methodology

The methodology consists of: (1) assembling and cleaning granular trajectory data, provided by Nokia HERE, (2) assigning these data to specific nodes that are already labeled with traffic control, (3) converting the trajectory data for each node into contour heatmap images, and



Figure 3.1: All nodes in the bounding box of the Nokia HERE data

(4) using these image data as both training and test data in the convolutional neural network, leveraging advanced computer vision models.

## Trajectory Data

In collaboration with the Connected Corridors project at UC Berkeley, we acquired data from Nokia HERE, which contain trajectories - speed, measured in meters per second, over time samples - from specific devices across the northeastern Los Angeles region, which includes Los Angeles, Glendale, Pasadena, Arcadia, and Monrovia. These devices include but are not limited to smartphones, connected cars and trucks, and loop detectors embedded in the roads. There are 19 million total data points, which is the product of the number of specific devices and each trajectory that the device captures. These data will subsequently be known as probe data.

## Street Network

We then obtain the street network of this region, based on the bounding box of latitude/longitude pairs of the data, from OpenStreetMap (OSM). This map is converted to a Python graph using the packages *OSMnx* and *geopandas*. The resulting graph contains roads (edges) and intersections (nodes) that are located within the area in question, as shown in Figure 3.1.

OSM, as an open-source project, is constantly updated by volunteers of the global community, and one major component of recent update has been specific traffic control at the network nodes. While far from comprehensive, the Los Angeles region fortuitously has sufficient traffic control information for the purpose of the project. In particular, there is a large number of nodes within the area in question that are labeled with a specific type of traffic control by OSM, within the choice set: traffic signal, stop, turning circle, or a motorway junction (exit/off-ramp). We then map the nodes that have labeled traffic control, as shown in Figure 3.2.




Figure 3.2: All nodes with traffic control in the bounding box of the Nokia HERE data




Figure 3.3: Sample trajectory in space from a particular device

## Data Processing and Cleaning

The first round of data cleaning involved dropping unnecessary columns and filtering data to specific providers with complete devices. There are both fleet and consumer devices, but this study narrowed down to several specific sets of consumer devices, namely CONSUMER21, CONSUMER14, and CONSUMER11, due to the high number of data points, the granularity, and the greater randomness in standard human movement, rather than relatively deterministic truck trajectories. We then convert the speed units and remove any missing values. In total, there are 19M points for the concatenated dataset of all three consumer devices.

The second round of data cleaning involves sorting the data by *here id*, a term from Nokia, of the data for each device. We then need to make sure each sample from each trajectory is occurring every 10 seconds at the most. If samples are greater than 10 seconds, then that point becomes the start of a new trajectory. Then, all trajectories with more than two data points are kept, while those with fewer are filtered out, since they are not useful if there is only one data point. Within the fully concatenated dataset, there are 63228 unique here ids. A sample trajectory is shown in Figure 3.3, which shows a particular device’s movement over space.

The next round of data cleaning is to find the nearest node with ground truth traffic control, previously shown in Figure 3.2, to each one of the trajectory points. This is done by calculating the Haversine distance between the points, which is in radians, shown

in Equation (3.14):

$$d = 2r \arcsin \sqrt{\sin^2 \frac{\phi_2 - \phi_1}{2} + \cos \phi_1 \cos \phi_2 \sin^2 \frac{\lambda_2 - \lambda_1}{2}} \quad (3.14)$$

where  $d$  is the distance between the two points,  $r$  is the radius of the Earth (specified as  $6.371 * 10^6 m$ ),  $\phi_1$  and  $\phi_2$  are the latitudes of points 1 and 2, respectively, and  $\lambda_1$  and  $\lambda_2$  are the longitudes of points 1 and 2, respectively. The Haversine distance is commonly used in navigation and geography due to its accuracy over a oblate spheroid surface like the Earth.

After determining the distances from each trajectory point to each node in the bounding box, a BallTree query is done to find the nearest node within that tree. The BallTree is used because of its efficiency of organizing points in a multidimensional space and ignoring points that are known to be further away than the point just explored. Since the units are in radians, we then convert back to degrees and then to a linear distance unit, specifically meters and feet. As a result, we now have 3501 unique nodes and 1899761 total points. Each labeled node has a varied number of nearest trajectory points, each of which contains speed and distance to the node.

The next step is to filter the nodes that have fewer than 100 points. The value of 100 is chosen arbitrarily, but intuitively we want to have a relatively dense number of points for each node without substantially decreasing the number of nodes available. This value can become a modifiable hyperparameter in future iterations of the project.

From here, we must now start visualizing the points for each node, specifically with respect to time, distance, and speed. In order to do this, a colormapped scatterplot is created, with the x-axis as time, y-axis as distance from the node, and the intensity of the color representing the speed, as shown in Figure 3.4. Note that the timestamps are in GMT, and thus it can be seen that much of the slower speeds occur at the morning and evening commute hours in PST. The nature of the heatmap does not require distinguishing time zones.

From this colormapped scatterplot, we must make this relatively sparse set of points denser in order for the image classifier to differentiate images more closely. The chosen method to do so is linear interpolation, which can fit curves using linear polynomials to construct new points within the scatterplot. After this linear interpolation for each node's data, we produce contour heatmaps, with an example shown in Figure 3.5.

## Neural Network

Our ultimate goal is to classify nodes accurately with their correct traffic control. This means our neural network will have 4 outputs, since we want to classify nodes into one of the four different classes noted above. The contour heatmaps produced for each node will be the input to the CNN. As specified above, each node's heatmap contains the time of each trajectory point, the distance from each trajectory point to the node, and the speeds of all trajectory points whose closest node is the specified node, and also linearly interpolated




Figure 3.4: Sample node with its associated trajectory points across time, distance from node, and speed.




Figure 3.5: Sample node with its associated trajectory points across time, distance from node, and speed, linearly interpolated to create a heatmap.

points with similar values. Each image is converted into a  $64 \times 64$  two-dimensional vector of pixels. This number was chosen arbitrarily, but may be adjusted to other powers of 2 for higher granularity. The color of the image must also be preserved within the pixelated array, as it is represented as a third-dimension with three possible RGB values. Finally, the traffic control label for each node is appended to each image for use in training the model, making it a supervised CNN. Within the 3501 nodes/images, 80% of them are randomly assigned to the training set and 20% are randomly assigned to the test set. This is a standard train/test split in the machine learning community, but can be altered as the designer sees fit.

From the *keras* library within TensorFlow, a Sequential model is used so that we can easily and linearly stack layers onto the neural network from input to output. Each image is of size  $64 \times 64 \times 3$ . In the initial design of the CNN, the output of the first layer is 32 channels, a standard specification. The kernel size is 5, which is both the height and the width of the 2D convolutional window. The stride controls how the filter convolves around the input, and is specified as 1 to reflect that the filter shifts one unit at a time. The activation function used in this layer is the ReLU, as specified in Equation (3.15).

$$f(x) = \max(0, x) \tag{3.15}$$

From here, a maxpool layer is applied. As mentioned before, the goal of this layer is to downsample the input to a pool size by a factor of 5. The hope is that the input layer produces a value that shows that there is a specific feature that exists in the image, and thus that feature's exact location is not as important as its location relative to other features. This can decrease the computational cost and also help to avoid overfitting.

The second convolutional layer follows the same skeleton as the first, but this has more output channels, with 50 (in the initial design). The reason for more layers is that the initial layer represents the high level features of the heatmap image and the subsequent layers represent more detailed features, which is typically reflected by more filters. The activation function is still the ReLU, as specified in Equation (3.15).

The third convolutional layer repeats the first and second, with the only change being an even greater number of filters, at 80.

After the three convolutional layers, there is one dropout layer, which is added to prevent overfitting. The dropout routine simply drops a random set of neurons by setting them to zero, with this dropping one quarter of neurons. This will cause redundancy in the system in such a way that the classification can still be made accurately without these neurons, which would be more specific to the actual input data. Note that we cannot do dropout on the convolutional layers because the same weights are reused, so it is not useful.

The output of the dropout layer is then flattened so that it can be fed into the dense, fully-connected layers. There are two layers, each having 512 neurons and ReLU activation. Each neuron will be a weighted sum of all of its inputs with the bias, which will then become the input to a nonlinear activation function.

After this, we do another dropout of one half. The final layer must output the number of possible classifications, which, in this case, is 4, where we want to label each intersection




Figure 3.6: Training accuracy using ReLU as the activation function for all layers except the output layer

as having a traffic light, stop sign, turning circle, or a motorway junction. The activation function in this last layer is the softmax.

The training itself will be driven with the cross-entropy as the loss function. Because there are more than two classifications (four), we use categorical cross entropy instead of binary cross-entropy.

## Results

The convolutional neural network described above trains the model and fits it to the training data with 100 epochs and a batch size of 100 in each epoch. The accuracy of the training data in the first five epochs hovers between 40% and 50%. The accuracy then gradually increases over each epoch, going up quickly between the fifth and tenth epochs and increasing at a pace between logarithmic and linear. It eventually converges to approximately 94%, as shown in Figure 3.6. The total number of trainable parameters, 5947064, reflects the total number of input weights and bias terms with every hidden input.

Accuracy and loss are inversely related, thus as the accuracy increases over each epoch, the loss decreases over that same time, starting at above 8.5 and eventually decreasing to .16 in the 100th epoch, as shown in Figure 3.7.

On the unknown test data, for which we have the labels but the model does not, it accurately classifies the traffic control at particular nodes 68% of the time. A sample classification and the success or failure is shown in Figure 3.10. This already is an improvement over random classification, as with random chance, it would be a 25% possibility of classification. It is also an improvement on using a uniform random classification, as the distribution of traffic control is 49% traffic lights, 31% motorway junction, 14% turning circle, 6% stop. With deciding simply to label every node as a traffic light, it would be 49% accurate. This shows an improvement of nearly 20% with the initial CNN.






Figure 3.7: Training loss using ReLU as the activation function for all layers except the output layer

An improvement to the model involves increasing the number of outputs for each convolutional layer. The goal of this is to increase the number of movement directions in the convolution itself, as a 3D matrix (or tensor) will require a 2D convolution that can convolve over more combinations of the vectorized image. By increasing each convolution layer from 32, 50, 80 to 50, 100, 200, respectively, the results improve. Specifically, the training accuracy increases to 99%, with the loss decreasing to .016, an order of magnitude lower than with the previous filter sizes, as shown in Figure 3.8 and Figure 3.9. This greater training accuracy and higher test accuracy suggests that the model still severely overfits on the training data, but still produces a higher test accuracy. Since a neuron in a CNN does a weighted sum across a small region of pixels, it intentionally operates on three dimensional chunks of data, with height, width, and the number of channels. When increasing the number of channels, it continues to reuse the same set of weights, slightly adjusted, as determined by the model from the image, which may explain the better test accuracy. We generally do not want a network that is not constrained, since then the neurons will directly remember the training data rather than having to make assumptions about features, such as vertical lines and colors, to emphasize when the images are more constrained.

In order to improve the model, many further hyperparameters can be adjusted, such as the activation function, the dropout rates, the type of optimizer used, and the number of neurons in the fully connected layers.

Instead of using the ReLU activation function, the function, as specified in Equation (3.16) was applied in all layers except the output layer. However, this produced a lower training accuracy of approximately 68% while also being much more computationally intensive.

$$\sigma(x) = \frac{1}{1 + e^{-x}} \tag{3.16}$$




Figure 3.8: Training accuracy with 50, 100, and 200 filters for the convolutional layers (with ReLU activation function)




Figure 3.9: Training loss with 50, 100, and 200 filters for the convolutional layers (with ReLU activation function)




Figure 3.10: A random set of 30 test images shows success rate of approximately 70% overall with ReLU activation.

The tanh function, specified in Equation (3.17) was also used in all the layers except the output, which under 40 epochs produced a 99% training accuracy.

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} \tag{3.17}$$

This produced the highest test accuracy at 72%, as shown in Figure 3.11 and Figure 3.12. The *tanh* function follows similar activation logic to the ReLU function, but because it contains several exponentials, it is considerably slower. Thus, the accuracy vs. efficiency tradeoff exists in this scenario, where the ReLU produces a comparable accuracy to the *tanh* function but in much shorter time.

Another hyperparameter that was adjusted is the learning rate itself. The learning rate can determine the magnitude by which the weights change at every epoch. A larger learning rate can lead to much greater noise and varying accuracies over time, while smaller training rates take much longer because the weights hardly change over each epoch. With the learning rate at .001, the training accuracy did not exceed 48%, reflecting the CNN’s inability to find weights due to the large discrepancies from previous weights. However, when the learning rate was decreased by an order of magnitude to .0001, the training accuracy skyrocketed, going above 90%. The learning is initially very slow, as shown in Figure 3.6, but after several epochs, the CNN begins increasing its performance very quickly.




Figure 3.11: Training accuracy with 50, 100, and 200 filters for the convolutional layers (with tanh activation function)




Figure 3.12: Training loss with 50, 100, and 200 filters for the convolutional layers (with tanh activation function)




Figure 3.13: A random set of 30 test images shows success rate of approximately 72% overall with tanh activation.

## Future Work

While this method produces significant improvement to the state of the art, the inference can still be made even more accurate. One method may be to use time lag trajectory points for the contour plot/heatmap. In other words, given the heatmap contains real data and interpolated data of the nearest trajectory points to each node, it may be useful to include the previous and following points in time of each point as well, regardless of whether those points are associated with the node by distance. This could weigh the temporal aspect of traffic control moreso than the distance.

Another approach is to use transfer learning. Image classification and computer vision are ripe areas for transfer learning because there are several large datasets, such as ImageNet, for which training models have already been developed. The filters in the first layer of the already developed CNN models may learn to recognize edges, gradients, and solid color areas; the second layer may have detected circles or angles; and so on. This can be also be used in a comparative study.

Yet another approach, as mentioned above, is to optimize more hyperparameters, perhaps by using a heterogeneous combination of loss and activation functions, or adjusting the number of convolutional layers. Countless combinations of hyperparameters can be done, but they must be considered in a reasonable and logical fashion.

As traffic control inference becomes more accurate, agent-based microsimulations, including MANTA, can become better models for human mobility. A similar methodology for

using image classification on trajectory data may unveil other insights as well, such as high betweenness centralities over time and space of particular nodes, which can be used for more optimal traffic engineering.

## Conclusions

The use of the state-of-the-art convolutional neural net (CNN) in order to infer traffic control of street intersections proves to be useful. While accuracy still can be improved, the results show that deep learning methods using computer vision and image classification of human mobility trajectory data is a promising technique for future inference problems. The tanh-based, three-layer CNN with 512 neurons produces the highest training and test accuracies, at 99% and 72%, respectively, showing that the model can be relatively generalizable without significant overfitting, classifying nearly 3 out of 4 nodes accurately on average. The ReLU-based CNN produces nearly comparable test accuracy at 70%, but is much faster because it does not do computations of exponentials, thus producing a designer tradeoff between accuracy and efficiency. This work can be the foundation of future traffic control inference that can be input into OpenStreetMap for various geographies across the world, given extensive trajectory data collected from smartphones, and eventually used in advanced microsimulation travel models.

Understanding and optimizing traffic control has been an open research area for decades. With the advent of autonomous vehicles, there is a hope that traffic control will be eliminated altogether, in favor of platooned and connected vehicles that negotiate intersections collaboratively. However, that reality remains decades away, and indeed the interaction between human-driven cars and autonomous vehicles will continue necessitating the understanding and simulating of traffic control.

## Additional ML Techniques in Simulation

Another important research area of urban simulation that can leverage advanced machine learning techniques is to improve the efficiency of macro- and micro traffic simulators themselves. Similar to using training data to determine the most accurate model for traffic flow prediction, it may be a significant improvement to the computational intensiveness of current simulators if it is possible simply to run only a few runs of a full simulator without losing accuracy in future runs. Specifically, with these few runs, each run's output data can be used as training to develop a nonparametric or deep neural network model of future simulator outputs. Indeed, this may allow the potential for much faster simulation, a major bottleneck in today's simulators [227, 209, 166]. One approach in doing so is to use the state-of-the-art traffic simulator, compute its benchmarks, train on only a few simulations, and then run the model on the outputs, producing metrics that may be validated against real simulator runs, ideally showing equal or better accuracies with larger speedups.

Yet another area is in leveraging deep reinforcement learning in agent-based models. Since much transportation simulation has moved toward activity-based and agent-based models,

deep reinforcement learning may be useful due to its ability to use exploration of solution spaces, or understanding how agents can develop specific behaviors based on a particular environment [164]. While agent-based modeling generally focuses on understanding behavior in an environment and reinforcement learning focuses on developing agents that can learn to negotiate that environment, determining how to use reinforcement learning rather than traditional logit and probit models in agent-level decision-making is an open area of research that can broaden the scope of urban simulation. Since reinforcement learning tackles the problem of correlating immediate actions with delayed returns, through an agent learning new information, it may be a natural complement, or perhaps even a substitute, to traditional discrete choice models [164].

### 3.3 Autonomous Vehicles (AV)

While machine learning continues to augment simulation and modeling, one new mobility paradigm, autonomous vehicles (AV), is also forcing the community to develop new capacity. Many questions remain about the impact of this new mode's integration into the transportation landscape, and simulations and modeling must consider these to inform policy and infrastructure.

AV is still in the inchoate stages of development and deployment. Major industry players, including Waymo, Tesla, Cruise, and Aurora have been progressing with near-market-ready AVs, as much of the enabling technology was developed in academia nearly two decades before [208]. Thanks to more powerful real-time computing, granular perception sensors and algorithms, advanced machine learning and computer vision, and ubiquitous connectivity, AV (and connected AV [CAV]) manufacturing is growing rapidly [17, 139].

Just as AV technology advances, simulation and modeling for AV has also been forced to evolve. The congestion effects of AV on cities have been found using agent-based microsimulations in different geographies [32, 80]. Within these microsimulations, traffic theory is now being amended to accommodate AV's capabilities. Since stop-and-go waves begin to occur after a particular threshold of vehicular density with human drivers, as Stern et al. note, much research has been devoted to learning how AV will influence bulk traffic flow [197]. Anderson et al. believe that AVs can affect congestion positively by reducing traffic delay, decreasing the number of accidents, and enhancing vehicle throughput, but may affect congestion negatively by increasing VMT [7]. In addition, Bagloee et al. note that because AVs can be fine-tuned to a great deal of precision, they can cruise at higher speeds while maintaining lower space headways between vehicles [17]. Stern et al. find that more than 20 vehicles operating in a circle can create waves almost instantaneously, primarily due to the arbitrary stop-and-go of human drivers. However, by controlling the velocity of a single vehicle in the flow, such as through injecting an AV into the circle, the waves begin to dissipate, reducing the standard deviation in the speed of all the cars in the circle by 50% [197]. This can have benefits not only to congestion, but also to fuel consumption, as the savings is as much as 40 percent when aggregated across all the cars in the traffic flow [197]. These studies

assume only a basic theory of autonomous control, and much of the technical development occurring is far more advanced than this system, reflecting even greater potential.

Similarly, Wu et al. developed Flow, a computational software platform that enables the development of reliable traffic controllers, including mixed-autonomy traffic in the aforementioned circle [232]. Using the deep reinforcement learning method of policy gradients, they develop different scenarios such as platooning vehicles, which improves the average velocity of the human-driven vehicles [232]. Qian & Yang assess AV impact on traffic flow as well, proposing a multi-class traffic flow model that encompasses the car-following behavior of both human-driven vehicles and AVs [173]. They model this mixed flow by considering the interactions of different classes of traffic streams, based on such parameters as the density of vehicles [173]. These streams stay consistent with the flow physics, but are data-driven compared to the theoretical LWR model [173]. Their model can accurately capture shockwaves of integrating multiple classes, and ultimately can produce realistic congestion propagation and accurate travel time predictions [173].

### 3.4 Urban Air Mobility (UAM)

As AV attempts to transform travel on the surface, urban air mobility (UAM) is doing so above the surface. Instead of adjusting transportation supply on the ground to increase the capacity of roads or make driving less unproductive, academic and industry leaders are envisioning a world in which daily transportation can occur in low altitude skies as well [127, 77]. As previously described, AV is now deeply within the scope of simulation and modeling, but UAM has only cursorily been researched in the simulation and modeling community.

In order to simulate UAM holistically, an integrated land use, travel demand, and transportation model would be a gold standard for city and regional planners [237, 225]. Basic templates exist from AV for understanding emergent modes' effects on cities, but they still must be extended [178, 229]. Specifically, simulations must develop the capability to do sensitivity analyses based on UAM adoption rates, network designs, and different variants of urban air transit and air taxi to find optimal deployments based on a set of constraints, such as minimizing total travel times or emissions [23, 237]. While aggregate approaches have been traditionally used in travel demand modeling, disaggregate activity-based models are also gaining traction and may be useful in this goal [99, 227]. Since such decisions as mode choice are conditioned on residential location choice and employment location choice, the daily activity pattern of an individual or household can be generated and thus be used as a more realistic input into the simulation models [175, 227]. Initial research has leveraged mesoscopic models and MATSim to evaluate UAM networks in Sioux Falls and Munich, but the contributions have been primarily in software integration and not in impact assessments [178, 180].



## Understanding the Potential Impacts of UAM

This section reviews potential impacts, considerations, and unanswered questions with respect to five thematic areas: (1) travel behavior/land use, (2) equity, (3) safety, (4) noise, (5) sustainability

### Travel Behavior/Land Use

Several analyses have attempted to estimate the potential impacts of UAM passenger services on travel behavior and land use. Many of these studies are based on four core assumptions. First, most of these studies assume that UAM will likely be a shared platform due to the high cost of aircraft [220]. Second, these studies generally assume that users will require smartphones for access, similar to transportation network companies (TNCs) [107]. Third, these studies generally anticipate that vertical take-off and landing aircraft (VTOL) will takeoff and land from vertiports, take-off and landing infrastructure for UAM in metropolitan areas that operate similar to helipads and heliports [77, 127]. Fourth, many of these studies assume that early UAM operations will start using an "air metro" operational model that includes aircraft flying along several initial high-volume travel corridors between a small network of vertiports. [216].

Under these assumptions, UAM's effect on travel behavior and urban form can be detailed more comprehensively. Emerging mobility technologies such as autonomous vehicles (AV) and UAM seek to tackle one of the known issues of mode choice and third party logistics operations, travel time, a theoretically and empirically known 'bad', or a factor that increases disutility [199, 89]. Both AV and UAM tackle this problem in different ways. UAM seeks to decrease *real* journey time, while AV seeks to decrease *perceived* travel time by reallocating time previously spent driving with more productive uses, such as working or sleeping. Both decrease generalized travel cost, but UAM's approach is arguably more direct. The anticipated decrease in generalized travel cost due to UAM may change travel behavior in several ways.

First, total VMT, either on the surface or in the air, may decrease, as some long distance travel or movement of goods over 30 miles could be shifted to aerial modes [220]. Circuitry factor, the ratio of distance traveled over the ground versus the distance traveled via air for the same OD pair, consistently exceeds 1.2 in the U.S. and 2 in other countries across the world [117], suggesting that miles traveled will almost certainly decrease. However, this is contingent on three issues. One is the efficient placement of vertiports to not only maximize the number of travelers, logistics, or emergency providers able to use UAM, but also to decrease the first mile/last mile connections to the vertiports themselves.

As described above, if the generalized cost of travel decreases in any form, then the bid-rent model suggests that there will be movement away from the central business district (CBD), which enables sprawl [6]. However, if the expectation in the U.S. is 500M UAS deliveries and 750M passenger trips by 2030, infrastructure devoted to the automobile, primarily parking and road capacity, may decrease [216]. Instead, urban mixed-use infill

development could follow transit-oriented development near vertiport stations, public transit stations, and along active and public transportation corridors. This may motivate more residential clustering by increasing accessibility of these highly attractive areas of high land value [104]. Lower density suburbs may also do infill development with the goal of creating higher accessibility areas near their local vertiports [70, 74, 35]. As a result, both centrality and dispersion may result. Specifically, the metropolitan urban form will likely be more polycentric and sprawled from a macroscopic scale, but each center may have higher densities and accessibilities near their local transit stations and vertiports [2]. Known issues with macroscopic sprawl include higher VMT and greater emissions from automobiles [33]. The sprawl that may result may also have an adverse impact on the environment if the aircraft are not fully electric, similar to AV [8, 73].

### Equity

A number of equity concerns have been raised with respect to UAM, most with respect to affordability and scalability. As it stands, many are concerned that UAM as a passenger use case will only be accessible to affluent households [140, 64, 100]. On-demand passenger services using helicopters in cities such as New York City, Mexico City, San Francisco, and São Paulo are premium products serving a predominantly wealthy clientele [224]. While electrification, aircraft autonomy, and scaled operations have the potential to increase usage among a broader market demographic, the extent to which these services could be affordable to a mass market is not well understood beyond early market studies based on a variety of assumptions [216]. Even with the most optimistic scenarios, mass market adoption may not be achievable, as UAM could still be unaffordable and inaccessible for disadvantaged communities, low-income households, and older adults.

UAM has two additional potential impacts on equity as well. First, as noted in the previous section, to the extent that UAM could reduce travel on the surface transportation network, UAM also has the potential to induce travel demand. If this were to occur, low-income travelers dependent on driving and public transportation would likely bear a disproportionate share of the negative impacts associated with induced demand, primarily increased congestion and longer travel times [96, 136]. Second, UAM could also directly compete with and reduce demand for commuter rail or long-distance public transportation, leading to increased fares and/or reductions in service. These changes may also disproportionately affect the underprivileged and minority groups, who are the highest proportion of public transit riders in many cities [120, 50].

Additionally, the placement of vertiports may also adversely affect vulnerable populations in a number of ways. Vertiports may be located in low-income and disadvantaged communities, causing unwanted noise, traffic (air and ground access/egress), and other environmental impacts [96, 48, 220]. Conversely, vertiports could be located in areas far from low-income centers due to lower demand, thereby disincentivizing usage by low-income households, creating an inequitable feedback loop.

Equity remains a notable challenge to UAM deployment, and UAM's effects on it and plans to address it vary considerably. Affordability, scalability, and infrastructure placement all may have detrimental effects on disadvantaged communities. Specifically, if UAM remains a niche service for goods or for passenger transport, these communities may never experience its benefits, and worse, may experience the negative externalities associated with UAM, including increased noise, congestion if vertiports are placed nearby, and emissions if initial VTOL aircraft are not electric [148, 155, 53]. If UAM is not able to scale considerably or if vertiports are placed further away from disadvantaged communities, rendering it effectively unusable by them, public resources used to create UAM infrastructure may have been unnecessarily diverted from mobility programs to help disadvantaged communities further, such as enhanced bus frequency and reliability [154]. Conversely, if vertiport location is accessible, costs for UAM decrease to affordable levels, and noise and emissions from aircraft are minimized or eliminated, UAM may have positive equity impacts. While vertiport location choice under equity constraints remains an open area of research, aircraft development to decrease noise and eliminate emissions while also bringing down cost is a priority for the nearly 250 companies currently operating in the UAM space [3].

## Safety

Safety represents a notable challenge for the deployment of urban air mobility. Both airspace users and the overflow are concerned about aircraft, and planners and policymakers must establish legal and regulatory frameworks to minimize any safety risks [220, 242, 77, 2]. Several studies have shown that the public typically associates UAM with helicopters, small aircraft, and general aviation, which have relatively poor safety records [147]. From 2003 to 2013, general aviation (GA) recorded 7.5 fatal accidents per 100 million miles traveled compared to 1.3 and .068 for the automobile and commercial aviation, respectively [148]. Due to the perceived association with general aviation and helicopters, UAM aircraft must be designed to overcome the perception of safety risk in addition to having equivalent or better levels of safety (EBLS) of commercial aviation, especially given their operation above more populated, thus riskier, environments [155, 220, 156, 148]. With aviation, the public tends to focus on fatalities per accident instead of total fatalities per year or relative safety to other modes [65, 152]. Due to association, UAM requirements will likely need to overcome this perception. To improve community and user perception of safety, successful small-scale demonstrations are needed before mass market deployment [140, 187].

Several significant *real* hurdles to an immaculate safety record include human error, inclement weather, and the lack of accurate sense-and-avoid technology in the air [73, 176]. In general aviation, 80% of accidents are caused by pilot error [148]. Thus, in manned UAM, [155] notes that pilots must maintain safe separation, vehicle control, and mission and vehicle management in order to minimize safety risks. In both manned and unmanned UAM, unpredictable wind gusts can cause the feedback control systems either not to sense the force quickly enough or to over-correct once sensed, causing unforeseeable behavior [169]. In unmanned UAM, sense-and-avoid technology is primarily driven by

advanced deep learning, sensing, and real-time visual data collection techniques, which are still in its inchoate stages [155]. In addition, *perceived* hurdles also exist, particularly in unmanned, autonomous aircraft, which communities fear will fall out of the sky [242].

Two different perspectives of safety exist, the rider in UAM aircraft and the overflown. Riders must be assured of their safety onboard the aircraft, likely at the order of magnitude of safety of commercial aviation. The overflown must be assured of the aircraft's stability and low probability of being in an accident above their homes [90, 242]. Legal and regulatory frameworks are currently being developed at different levels of decisionmaking, including unmanned traffic management in the airspace by NASA, and both pilot and drone certification by the FAA and EASA in Europe [194, 20, 148, 156].

### Noise

Another major constraint for UAM will be noise in the skies. As UAM aircraft will be multicopter VTOL, they are particularly susceptible to noise, which is caused by each blade entering the region of the previous blade while rotating [8, 39, 82]. Two scenarios may unfold with UAM: (1) a single, infrequently flown aircraft with an irritant noise that subsides as it quickly travels away, and/or (2) constant noise from many aircraft flying in the skies for many hours of the day. Decreasing this sound exposure level (SEL) over time for either scenario is currently a notable challenge in both aircraft and rotor design as well as planning of facility location and airspace restrictions [215]. For goods mobility on smaller aircraft, experts have targeted SEL below 60 dB at heights up to 300 feet [186].

However, it is not only the volume that matters. [241] found in Los Angeles, Mexico City, Switzerland, and New Zealand that the type of noise (more mathematically known as frequency), in addition to volume, is also a concern for communities being overflown [241]. The known reference of acceptable sound is that of a bee buzzing, which respondents have found to be less obtrusive in both volume and type of sound [241]. However, communities vary in their preferences [88, 241].

Flight altitudes remain a key variable to SEL. The higher the altitude, the lower the SEL, but higher altitudes above 10,000 feet also drain batteries more quickly on takeoff and landing functions, thus constraining UAM aircraft to lower airspace. In addition, higher altitude flying may also interfere with commercial aviation in class A, which is not acceptable [206]. As a result, class B, C, and D airspaces around major airports and above communities are expected to be used, but these are at lower altitude between 300 and 3000 feet, making them much noisier to those on the ground [206]. Because of noise concerns, many communities have restricted helicopter routes, resulting in fragmented local, state, and federal regulations and guidance. Aircraft manufacturers and operators may need to adapt to this existing ecosystem of regulations in order for communities to integrate UAM in low-level airspace [216, 206].

Accordingly, the diverse legal and regulatory frameworks across different communities will also prove challenging from both the technical and planning lenses. Theoretically, legislation and regulation could enable communities to establish noise limitations when airspace users

develop flight plans (a concept known as constrained flight plans) [127]. If UAM flight paths must avoid noise sensitive land uses (versus taking the most direct route), operations could become increasingly complex, creating challenges associated with battery range and travel time savings. Additionally, different communities, even if adjacent to each other, have different thresholds for noise, making constrained flight plans particularly challenging [222].

## Sustainability

As multimodal mobility becomes more prevalent, the sustainability impacts of each mode need to be objectively evaluated to preempt future deployment. In the U.S., transportation accounts for 28% of all greenhouse gas (GHG) emissions, the highest proportion of any sector. Within the transportation sector, 83% of GHG emissions are emitted by automobiles and trucks [217]. While sustainability research in EVs and micromobility continues, UAM's remains unclear and its research agenda more complex [faria, 144]. Three major questions must be asked with respect to sustainability for UAM: (1) Do UAM aircraft itself generate more emissions? (2) Will UAM create more sprawl, and if so, will that increase lead to more emissions? (3) Are there indirect operations and outcomes of UAM that will lead to more emissions, such as deadheading? In this section, these three questions will be addressed.

**Aircraft Emissions** Many manufacturers have a goal of certifying electric vertical take-off and landing (eVTOL) in the 2020s. These aircraft could be zero emission, if powered by a clean electric grid; however, life cycle emissions are often not considered as part of this calculus. Several aircraft manufacturers have begun exploring hybrid-powered aircraft designs as well, primarily to enable power intensive operations such as hovering [147, 117]. Internal combustion engine (ICE) VTOL is also being explored for particular use cases that may be one-off or naturally have lower environmental footprints [2].

Several future scenarios of aircraft distribution exist.

1. eVTOL
2. eVTOL, hybrid VTOL
3. eVTOL, hybrid VTOL, ICE VTOL

The extent of UAM aircraft emissions not only depends on the likelihood of these scenarios, but also the share of each mode in each scenario. Both the scenario likelihoods and mode shares are heavily dependent on expected use cases. Case (3) may be the most likely because it can encompass a variety of use cases in both passenger and goods UAM. For example, [2] found that ICE VTOL are more appropriate for operations that are more infrequent but require longer flight times and range, such as rural emergency response [2]. For the use case of passenger UAM, eVTOLs emit fewer GHGs on a VKT basis compared with ICEVs for trips beyond 35 km, but the average ground-based vehicle commute is only 17 km, and trips exceeding 35 km account for only 15% of all vehicle trips, suggesting that there is a niche passenger service that may use eVTOL or hybrid VTOL. [117]. However, while one 100 km

point-to-point trip with one pilot results in 35% lower GHG compared to an ICEV, it is also 28% higher compared to an EV [117]. Indeed, if EVs significantly penetrate the market by the time UAM is deployed at scale, UAM's environmental benefit will likely decrease. If UAM's life cycle emissions are less than EV's life cycle emissions, there will be benefit to widespread UAM deployment for both passenger and goods.

**Sprawl** Location theory in urban economics suggests that as transport costs are reduced, movement away from the center occurs [96]. Thus, if passenger UAM is adopted at scale for longer distance trips, it is likely that the urban boundary expands as citizens and employers move further away, contributing to polycentric urban sprawl [5, 192, 43]. This sprawl may lead to even greater transport costs and potential emissions across the suburban centers, particularly if freight and passenger vehicles do not move toward cleaner fuel sources [217]. However, if light freight and cargo delivery use goods UAM within smaller areas to replace inefficient routes of last-mile delivery trucks, there can be sustainability benefits even with the added sprawl [40] Alternatively, on a local scale, accessibility may increase near potential vertiports due to more urban infill and more activities that are conducive to active modes of transit, also positively contributing to sustainability [24].

**Indirect effects** Depending on the scenario and use case for UAM, several other impacts may occur. If passenger UAM is deployed at scale and effectively replaces long-distance drivers, the new perceived increase in level of service on the road may induce further demand for the car and subsequently increase congestion and emissions altogether [76, 122]. However, the scale to which UAM must penetrate to realize an induced demand effect remains unknown.

Passenger UAM may also be affected by ongoing and future pandemics, such as COVID-19. With COVID-19 forcing many employees to work remotely, air pollution and emissions have decreased across the world [3]. As of this writing, it is unclear to what degree future work will be done remotely [3]. If this proportion is high, congestion may decrease altogether as there is less travel demand. However, if the proportion of remote work remains low, then the level of road congestion may increase if the metropolitan area heavily depends on public transit, which may then push travelers to UAM [109]. The effect of this change will vary depending on the type of aircraft.

In addition, the exact paradigm of passenger UAM will dictate various impacts. For example, if the aerial taxi paradigm prevails, then deadheading, in which aircraft continuously fly around until a passenger requests service, will become significant. Studies have shown that deadheading in TNCs, in both commuting and between rides, account for nearly 45% of all TNC VMT, and thus emissions [230]. This same behavior may occur with passenger UAM since parking the aircraft in specific places will not be optimal for routing or scheduling [77]. If it is an urban air metro paradigm, the capacity of the aircraft, the headways between them, and the transfer delays at vertiports will dictate its ability to substitute travelers currently using less green modes [216].

Meanwhile, goods UAM may also find a niche due to ongoing and future pandemics. Household consumption of delivery services has increased substantially [18]. If these deliver-

ies are currently using ICEVs, then their replacement with EVs or goods UAM will become a sustainability improvement in the future, depending on the degree of adoption [18].

## Contributions

Chapter 5 and Chapter 6 of this dissertation provide novel simulation evaluation tools and analytics that help to understand the impacts of UAM on travel behavior. Indeed, the state of the art to understand UAM's impact on the urban environment - be it in sustainability, equity, or efficiency - is still young, and the forthcoming contributions presented in this dissertation attempt to further the research agenda as the excitement and interest in this new mode continues [28]. Chapter 7 focuses on urban air mobility and its introduction into cities, specifically whether local and regional government agencies are prepared for this new mobility mode and the impacts to their communities that should be expected if deployed at scale.

# Chapter 4

## Planning and Designing Land Use-Constrained Networks of Urban Air Mobility Infrastructure

### 4.1 Introduction

Urban air mobility (UAM) is a concept that envisions safe, sustainable, affordable, and accessible air transportation for passenger mobility, cargo delivery, and emergency management within or traversing a metropolitan area [127, 77, 167, 19, 54]. In recent years, developments in electric and hydrogen propulsion; vertical and short takeoff and landing (VTOL/STOL) aircraft; and other technologies are creating renewed interest in on-demand urban aviation [127, 77].

While UAM may be enabled by the convergence of many factors, several challenges, including community acceptance, safety, equity, planning and implementation, airspace management, and operations, could create barriers to mainstreaming [54]. The deployment of UAM will require a network of landing facilities; charging and/or refueling infrastructure; and digital infrastructure for advanced air traffic management. While UAM may be able to use existing helipads initially, growth of the sector will likely require additional facilities for aircraft (e.g., takeoff and landing, parking, and maintenance) and passengers (e.g. waiting, boarding, access, and egress) [77, 107, 150, 180]. However, establishing new facilities could face a variety of barriers, such as local opposition, infrastructure costs, and multimodal integration [241, 147].

Several market studies conducted prior to the global pandemic estimate a market potential for UAM passenger service between \$2.8 to \$4 billion US in the United States by 2030 [142, 68]. However, Johnston et al. (2020) and Lineberger et al. (2019) find that limited infrastructure has the potential to constrain the growth of UAM [68, 142].

A number of studies have attempted to develop standard terms, definitions, and a taxonomy of takeoff and landing infrastructure sizes and configurations [142, 68]. Generally,



these studies classify three sizes of UAM takeoff and landing infrastructure, from largest to smallest: 1) vertihubs; 2) vertiports/vertibases; 3) and vertipads/vertistations. Table 4.1, adapted from Cohen et al., provides additional information on these facilities [54]. Other terms include skyports and aerodromes. UAM may also be able to use existing infrastructure such as heliports, airports, and waterways (for amphibious operations). In this paper, the authors will refer to UAM infrastructure as vertiports [21].

Table 4.1: Taxonomy and Definitions of UAM Infrastructure (adapted from Cohen et al.)

Type of UAM infrastructure	Description
Vertihub	A larger facility (possibly with multiple floors) to accommodate numerous landing pads with parking for multiple aircraft. Johnston et al. (2020) estimate vertihubs would be approximately 70,000 square feet, spread across multiple floors, and cost between \$6 and \$7 million US to build.
Vertiport/vertibase	A medium-sized facility intended to accommodate up to three landing pads and up to six parked aircraft. Johnston et al. (2020) estimates vertiports/vertibases would be approximately 23,000 square feet and cost between \$500,000 and \$800,000 US to construct.
Vertipad/vertistation	A single landing pad and parking stall intended to accommodate one or two parked aircraft. Johnston et al. (2020) estimate that vertipads would be approximately 6,000 square feet and cost between \$200,000 and \$400,000 US to construct.

In a fully mature network, Johnston et al. (2020) estimate that large cities (e.g., London, New York City, and Tokyo) could require 85 to 100 takeoff and landing pads comprised of three to five vertihubs, ten to fifteen vertiports, and five to ten vertipads [142]. The study estimates that this infrastructure would require a capital investment of \$35 to \$45 million US with annual operating costs of \$110 to \$130 million US per year. In medium-sized cities (e.g., Atlanta, Denver, and Düsseldorf), the study estimates a network of 38 to 65 takeoff and landing pads spread across two to three vertihubs, five to ten vertiports, and three to five vertipads. The investment required for medium-sized cities is estimated at \$15 to \$20 million US for capital costs and \$35 to \$50 million US annually for operations. Johnston et al. (2020) estimate that electric charging infrastructure could represent up to 75 percent of the total

initial capital costs and that electricity could also represent 30 to 35 percent of the estimated annual operating expenses [142]. The study concludes that UAM infrastructure will likely be smaller scale during initial service deployment and will require a combination of private sector investment, public sector subsidies, and non-aeronautical revenue sources (similar to revenue generated by airports from retail, landside access, and other sources) [142].

## Vertiport Siting and Network Design

Several studies have examined vertiport siting and network design both qualitatively and quantitatively. Fadhil et al. (2018) utilizes a weighted linear combination (WLC) method to identify potential vertiport locations in both Los Angeles and Munich using qualitative data from subject matter experts [79]. Rothfeld et al. (2018) uses a similar weighted combination based on qualitative location suitability and a location-allocation optimization algorithm [179, 178, 180]. The study uses MATSim, an agent-based simulation framework for travel behavior analysis. Daskilewicz et al. (2018) analyzes the spatial distribution of jobs, income, and population in the San Francisco Bay Area and Los Angeles to maximize travel time savings [62]. The study concludes that optimal vertiport locations would likely be in higher income census tracts.

Uber Elevate (2016) develops a k-means clustering method using long-distance Uber rides to identify potential vertiport locations in Los Angeles and London [77]. However, because the study uses proprietary data from one service provider, the results are limited. Bulusu et al. (2020) also uses a k-means clustering model to explore the San Francisco Bay Area [42]. However, demand data for the study are not comprehensive for the entire metropolitan region, an important consideration to understand the broader impacts across the transportation network [91]. Lim & Hwang (2019) and Syed et al. (2017) also use k-means clustering in Seoul and Washington D.C., respectively [132, 203]. Both studies find that the location of vertiports is more important than the number of vertiports. One notable downside to k-means clustering is that vertiport locations are susceptible to outliers, which may place vertiports further away from the majority of demand. If vertiports are further away, then demand for UAM may change altogether. Additionally, some vertiport locations considered suitable in the model may be impractical in the real world when considering noise, equity, and other concerns.

A number of studies have found that UAM adoption could be influenced by vertiport placement [19, 178, 188]. This study suggests feasibility constraints, such as zoning and parcel areas of potential locations, should also be considered. Robinson et al. (2018) uses vacant land parcel information in South Florida to understand the potential for short takeoff and landing aircraft. However, the larger parcels required for STOL operations may not be needed for VTOL aircraft [177]. Wu & Zhang (2020) develops a single-allocation hub-and-spoke optimization formulation for Tampa, Florida. The study uses LiDAR and Geographic Information Systems (GIS) to identify potential vertiport locations based on available ground spaces and land use restrictions [233].




Figure 4.1: The San Francisco Bay Area Network

Antcliff et al. (2016) suggests the San Francisco (SF) Bay Area could be an early adopter market for UAM because of the high percentage of long-distance or “super” commuters, weather, and geography [8]. In addition, the existence of multiple urban centers in the region tends to result in bidirectional traffic patterns that could reduce “deadhead” flights without paying passengers and improve the operational efficiency and profitability of early operations [193].

This paper is organized into three sections. The first section introduces a novel modeling methodology that could be used to estimate the number of recommended vertiports, using fast, high-fidelity regional-scale microsimulation. The second section highlights the results of the model. The final section concludes with a discussion of study limitations and proposes recommendations to help advance UAM infrastructure planning, implementation, and community integration.

## 4.2 Methodology

Vertiport network design requires determining the location of every vertiport in the metropolitan area with all the other locations of vertiports in mind. As such, we introduce a clustering algorithm for optimal vertiport locations, based on existing travel demand in the SF Bay Area. These locations are specifically made to be feasible by introducing land use and zoning constraints. In this section, the microsimulator, street network, demand, network optimization, land use constraints, and overarching simulation procedure are described.

## Microsimulator

This paper uses Microsimulation Analysis for Network Traffic Assignment (MANTA), a fast, parallelized, GPU-based regional-scale traffic microsimulator. MANTA can simulate millions of trips along a given street network in only a few minutes, using the intelligent driver model, gap acceptance, and lane changing algorithms to move the vehicles on each edge as they interact with other vehicles [239]. This version of MANTA also leverages a priority queue Dijkstra’s algorithm for the routing of trips. MANTA has been calibrated and validated with open source Uber Movement data [239].

MANTA is used over existing simulators for several reasons. In transportation simulation, tradeoffs exist across granularity, spatial scale, computational time, and simulator costs. While many simulators exist that satisfy several of these, none is able to satisfy all four to enable quick scenario planning. A high granularity is necessary for UAM analysis because of the sensitivity of UAM demand to variations in travel time [77]. The spatial scale also needs to encompass greater metropolitan regions since UAM is expected to address longer trips by distance and time [77, 19]. Due to the high granularity and large spatial scales necessary for analysis, computational times of existing simulators are too long – on the order of days – for fast sensitivity analysis [239]. Since MANTA is a GPU-based simulator, the cost of the processors is also more affordable than more powerful machinery such as high-performance computing (HPC) architectures, which allows it to be accessible to researchers [239].

## Network

The street network is constructed from the OpenStreetMap (OSM) network of the San Francisco Bay Area using the OSMnx library. The network is cleaned to remove nodes at curves or bends in edges in OSM that are not topologically accurate [34]. The resulting network contains 224,223 nodes and 549,008 edges. The metropolitan region’s nodes are shown in Figure 4.1.

## Demand

This study uses the San Francisco Bay Area’s Metropolitan Transportation Commission’s (MTC) Travel Demand Model One. The MTC model uses a population synthesizer that locates households based on the 2000 Decennial Census Public Micro-sample, where each traffic analysis zone (TAZ) typically has a similar population and can range from the size of a block to several square miles [143]. It is then validated across years 2000, 2005, 2010, and 2015. Using this behavioral data, this study uses a granular microsimulation that randomly assigns origins and destination, sampled from a uniform distribution, within a traffic analysis zone to a network node in that TAZ [226].

A simplified method is used to determine demand for UAM, by filtering automobile trips with travel times  $\geq 30$  minutes, known as Marchetti’s Constant [138]. Marchetti’s Constant is limited by its lack of true behavioral realism, but can be a useful tool because

of its simplicity [25]. A preliminary run in MANTA is performed using 3.2M trips from San Francisco Bay Area MTC between 5 AM and 12 PM on a typical day, and the resulting travel times are then filtered to  $\geq 30$  minutes [239].

## Network Optimization

The vertiport network design problem can be viewed as an uncapacitated variant of the facility location problem (FLP), a common formulation in transportation and operations research [71]. This problem is uncapacitated because the maximum number of aircraft is not specified for each facility. The objective is to minimize the sum of the distances traveled to and from every traveler’s nearest vertiport, based on every traveler’s origin and destination. The uncapacitated FLP is shown in Equation (4.1).

$$\text{minimize} \quad \sum_{i \in V, j \in D} d_{ij} x_{ij} \quad (4.1a)$$

$$\text{subject to} \quad \sum_{i \in V} x_{ij} = 1 : \forall j \in D, \quad (4.1b)$$

$$\sum_{i \in V} y_i \leq k, \quad (4.1c)$$

$$\forall i \in V, \forall j \in D : x_{ij} \in \{0, 1\}, \quad (4.1d)$$

$$\forall i \in V : y_i \in \{0, 1\} \quad (4.1e)$$

where  $V$  is the number of vertiports,  $D$  is the number of travelers (demand),  $x_{ij}$  is whether traveler  $j$  is connected to vertiport  $i$ , and  $y_i$  is if the vertiport is open to use. The first constraint in Equation (4.1b) indicates that every demand point is assigned to only one vertiport. The second constraint in Equation (4.1c) requires that the number of active vertiports must be less than or equal to the specified  $k$ , the maximum number of vertiports. The third and fourth constraints in Equation (4.1d) and Equation (4.1e) indicate that  $x_{ij}$  and  $y_i$ , respectively, are binary variables.

In this study, this linear programming problem is reformulated using an unsupervised  $k$ -medians clustering approach. Thus, for each cluster, the objective function transforms into Algorithm 1. The first step of the algorithm is to randomly select  $k$  points from the demand as cluster centroids, sampled from a uniform distribution. Then, at every iteration  $t$ , each node is added to a cluster based on its minimum distance to each of the  $k$  centroids. The next step recalculates the median of that cluster, and every new median cluster point has its cluster points  $d$  recalculated, again based on the minimum distance of each demand point to the nearest of the  $k$  clusters. Finally, the algorithm converges to a set of cluster centroids with its associated demand points.

The  $k$ -medians algorithm calculates the L1 norm absolute (Manhattan) deviation between points [61]. This contrasts with the squared deviations computed in a  $k$ -means problem.

**Algorithm 1** k-medians Clustering for Vertiport Location

---

```

1: procedure K-MEDIANS( $D$ )
2:    $D = \{d_1, d_2, \dots, d_n\}$  is the set of input demand points
3:   Randomly initialize  $K$  clusters,  $C = \{\eta_1, \eta_2, \dots, \eta_k\}$  according to uniform distribution
4:   while  $|\eta_{k,t+1} - \eta_{k,t}| > \epsilon$  do
5:     for  $\{n = 1, \dots, N\}$  do
6:       Find center  $\eta_k \in C$  closest to point  $d_n$ 
7:       Assign demand point  $d_n$  to center  $\eta_k$ 
8:     for  $\{k = 1, \dots, K\}$  do
9:       Calculate median of all points in cluster  $k$ 
10:      Assign new center  $\eta_k$  as median

```

---

## Land Use Constraints and Parcel Assignment

As previously discussed, land use and zoning requirements can limit vertiport placement [226]. To plan a vertiport network, additional factors should be considered in addition to UAM demand. A holistic optimization must be executed to understand not only what is *efficient*, but also what is *possible*.

Using parcel level data from the San Francisco Bay Area, land use is constrained to parcels that are zoned office or mixed-use office space, light and heavy industrial (including warehouses), retail, and parking structures [147, 220, 77]. These parcels are believed to have compatible zoning for vertiports and minimum requirements that allow for take-off and landing operations at a feasible scale. The nearest vertiport parcels that satisfy these land use constraints are assigned using a fast BallTree nearest neighbor search. This technique recursively divides the data into nodes defined by a centroid to perform a single distance calculation for lower and upper bounds [159].

Additionally, the parcels that support larger vertiports are also considered. Cohen (1996) suggests that the area needed for helicopter takeoffs and landings should be  $\geq 14$  acres [55]. However, literature on eVTOL suggests that vertiports may require a considerably smaller footprint of approximately 1.4 acres (or 60,000 square feet) [81]. Because of the diversity of aircraft currently under development, the model considers parcels between 1.4 and 14 acres as potentially suitable for vertiport infrastructure [220].

## Simulation

After potential eligible vertiport parcels are identified, the model simulates how the parcel locations may impact travel behavior. The nearest node to each vertiport parcel is used as a proxy point for later input into the MANTA traffic microsimulator. Please refer to Figure 4.2 for an example. Trips that have the same origin and destination vertiports are eliminated and assumed to be a driving trip (e.g., the trip is too short to be served by UAM). The subsequent analysis assumes that the remaining eligible trips have been assigned as UAM




Figure 4.2: In San Francisco, a sample vertiport parcel (red) determined from the clustering algorithm and its nearest node (orange) that is used in subsequent microsimulation

trips.

The demand used in the original ground travel simulation run in MANTA is parsed into trips from (1) the origin to its nearest origin vertiport, (2) origin vertiport to destination vertiport, and (3) the destination vertiport to the destination. The travel time of the first leg from the origin to the nearest vertiport is computed with an initial run in MANTA. The travel time from vertiport to vertiport is calculated using the Haversine distance, which is the great-circle distance between two points on a sphere given their longitudes and latitudes. The aerial travel time is then calculated assuming a flight speed of 150 mph [77]. The travel time from the destination vertiport to the final destination is then simulated in MANTA. To calculate an accurate departure time for those trips requires both the previous legs' travel times and estimated transfer delays at the vertiports, as shown in Equation (4.2).

$$T_d = T_o + T_{o,over} + T_{over,dvert} + 2T_{vert} \quad (4.2)$$

where  $T_d$  is the departure time at the destination vertiport for the final destination,  $T_o$  is the departure time of the traveler from the origin to the origin vertiport,  $T_{o,over}$  is the travel time via automobile from the origin node to the origin vertiport,  $T_{over,dvert}$  is the aerial travel time between vertiports, and  $T_{vert}$  is the transfer delay at each vertiport.  $T_{vert}$  is assumed to be 10 minutes at each vertiport, based on the maximum acceptable waiting times for transit, security, embarkation, and disembarkation [10].

### 4.3 Modeling Results

This section details the results from both the clustering and the simulation processes. The clustering process determines the location of the vertiports from the k-medians algorithm, and then maps that parcel to the nearest feasible parcel – and subsequently, node – based on the land use constraints. The regional-scale simulation uses the MANTA traffic microsimulator to calculate the final congested travel times of each driving and multi-modal UAM trip.




Figure 4.3: Converged map of  $k = 10$  original cluster centroids across the San Francisco Bay Area, in radians, based on OD nodes of the street network

## Clustering

Due to the randomization of the initial points, several trials are run to determine the final convergence points for the  $k$ -medians clustering algorithm. Convergence occurs between 3 and 15 iterations, depending on  $k$ . At  $k = 10$ , the cluster points, in radians, are shown in Figure 4.3.

When  $k$  is low, the mean distance between the calculated cluster point and the nearest feasible vertiport parcel is high, but as  $k$  increases, the mean distance of parcels to the initially calculated median points stabilizes between 15 to 23 km, as shown in Figure 4.4. This large distance suggests that feasible vertiport sites may be challenging to find and thus limit the overall value of UAM altogether.

While a greater number of vertiports may reach more potential users within a shortest distance from a vertiport, many of the initial cluster points found are similar but simply farther from usable parcels. This suggests that land use patterns may be similar across localities. Resulting vertiport parcels are shown in Figure 4.5 for  $k = 10$  and  $k = 50$ .

In Figure 4.5,  $k = 10$  creates a sparse network while  $k = 50$  creates a much denser network of vertiports. When  $k$  increases, the vertiports appear in the major low-density suburbs adjacent to San Francisco Bay itself, often near a large thoroughfare or highway (e.g. CA-82 in the Peninsula, Interstate 880 in Oakland, and Interstate 280 in San Jose). Interestingly, San Francisco, with over 800,000 residents and a dense financial district, shows only two feasible parcels in the  $k = 50$  design and only one feasible parcel in the  $k = 10$  design. This suggests zoning constraints could have a notable impact on the potential number of UAM vertiports [162].

From these parcels and the calculated departure and travel times for UAM trips, the model simulates the UAM trips using MANTA.






Figure 4.4: Distances between calculated clusters and feasible parcels shows that as  $k$  increases, the mean distance tends to decrease. The green triangle represents the mean distance and the orange line represents the median distance. The box and edges show interquartile range, minimum, and maximum distances.

## Simulation

Several analytical assessments are made from the simulator outputs. The first explores average UAM travel times across  $k$  compared to non-UAM driving trips. The mean travel times of UAM and non-UAM across  $k$  vertiports are shown in Figure 4.7. UAM shows benefit over driving the same trip at  $k \geq 40$ , but shows benefit across both modes together in the transportation network only once  $k \geq 100$ .

The kernel density estimation of the time deltas between the UAM and non-UAM scenarios across  $k$  are shown in Figure 4.6. In this plot, a negative number indicates that a UAM trip's time is longer than the same trip's time when driving. A positive number indicates that driving takes longer than UAM for the same trip. The mean delta in travel time for  $k = 10$  is -32 minutes,  $k = 20$  is -13 minutes, and  $k = 30$  is -7 minutes, showing that at low  $k$ , the travel times from origin to destination are significantly faster when driving compared to using UAM. However, the mean delta in travel time becomes positive at  $k \geq 40$ , suggesting there must be at least 40 vertiports for the average UAM trip to be faster than the average driving trip. A significant decrease in UAM travel time - nearly 20 minutes - occurs from




Figure 4.5: Adjusted cluster centroids (vertiport parcels) across the San Francisco Bay Area that satisfy local land use requirements for  $k = 10$  (blue) and  $k = 50$  (orange)

$k = 10$  to  $k = 20$  and from  $k = 60$  to  $k = 70$  vertiports. At  $k \geq 70$ , the rate of change in travel time begins to plateau again, suggesting that  $k = 70$  may be a good compromise for the number of vertiports in the San Francisco Bay Area to provide travel time benefit.  $k = 100$  vertiports is an inflection point in which the mean travel time for trips with both UAM and driving becomes lower than the mean travel time for only driving across all trips.

Even using Marchetti's Constant, not all trips would logically be assigned to UAM. Figure 4.8, where  $k = 10$ , shows an example of a trip where UAM travel time is worse than driving. In this figure, the vertiports (red and light blue, respectively) are out of the way from the origin and destination (peach and green, respectively), increasing the total travel distance and time.

On the other hand, an example of a UAM trip with a better travel time than driving is shown in Figure 4.9, where  $k = 70$ . In this figure, the vertiports are along the route of a vehicle trip, reducing the travel time that would otherwise be required to deviate to the nearest vertiport.

The second assessment is to understand how area constraints of the parcels affect vertiport locations and subsequent travel times. Given  $k = 70$ , Figure 4.10 shows the mean travel times of UAM compared to the non-UAM scenario when adjusting for area of the parcel needed,  $A$ , from 2 to 14 acres by increments of 2. Intuitively, the mean travel times with UAM increase as the area constraints increase, confirming that land parcels of large areas such as 14 acres are less accessible and less prevalent. Only if the acreage constraint is

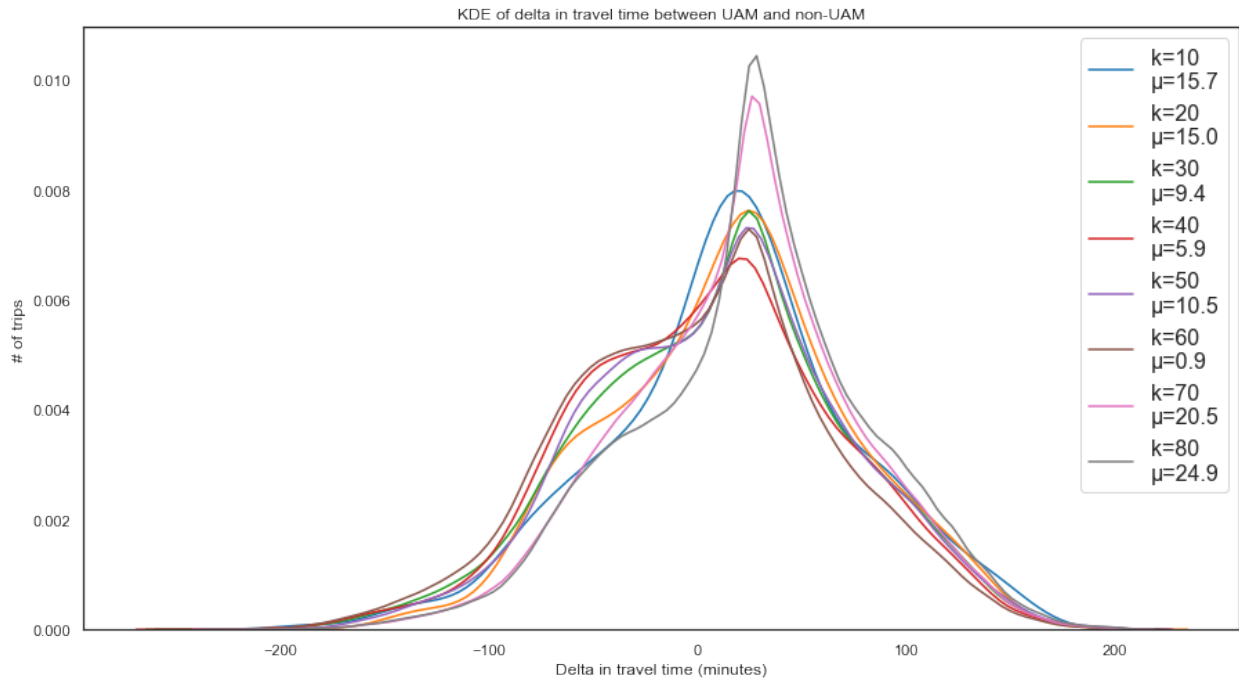


Figure 4.6: Difference in travel times between driving and UAM across  $k$  when parcels constrained to industry, retail, office, and parking garages, and to 2 acres or larger. Positive values imply longer non-UAM travel times and negative values imply longer UAM travel times

below 2 acres does the mean travel time of UAM remain competitive with driving. This result suggests that a larger network of smaller vertiports makes more sense than a smaller network of larger facilities, as described in Table 4.1. It may be necessary to build vertically or create seamlessly multiplexed FATOs and TOLAs in small spaces. With the expectation that UAM use VTOL instead of STOL aircraft, parcel areas  $A \leq 2$  acres are feasible if vertiports are designed efficiently [206, 130].

## Limitations and Future Work

Several limitations exist in the model. The first is the aerial travel time assumption, which currently is computed by a Haversine distance metric and a static assumption of aircraft speed at 150 mph. As of May 2021, there were over 200 aircraft under development, and the average speeds of these aircraft are unknown. The precise times of each component of the flight, such as takeoff, cruise, and landing are also not considered [117]. In addition, exact flight trajectories that consider collision avoidance and weather are not incorporated. Ongoing work includes the integration of accurate unmanned aerial simulation [235].




Figure 4.7: Mean travel times across  $k$  in UAM scenarios vs. mean travel time in a non-UAM scenario when parcels constrained to industry, retail, office, and parking garages and to 2 acres or larger. For  $10 \leq k \leq 30$ , UAM takes longer than driving for the same trip. For  $40 \leq k < 100$ , driving takes longer than UAM for the same trip.  $k \geq 100$  produces shorter mean travel times across all trips.




Figure 4.8: The origin, destination, and vertiports in an example UAM trip. Origin in peach, destination in green, vertiport nearest the origin in red, and vertiport nearest the destination in light blue. This trip would not benefit from UAM because the vertiports are far from the origin and destination and in opposite directions of the desired direction of travel. The driving distance between origin and destination is 15 miles, while the driving and aerial distances together with UAM is approximately 38 miles.

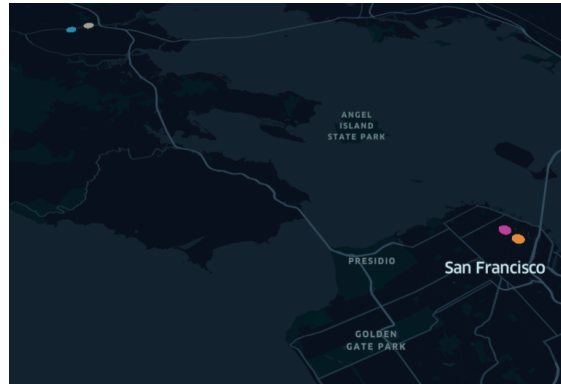


Figure 4.9: The origin, destination, and vertiports in an example UAM trip. Origin in blue, destination in orange, vertiport nearest the origin in brown, and vertiport nearest the destination in magenta. This trip benefits from UAM because the vertiports are along the way if the trip was made driving. The driving distance between origin and destination is 18.6 miles, while the driving and aerial distances together with UAM is approximately 15.5 miles.




Figure 4.10: Mean travel times for UAM and non-UAM scenarios as area constraints of vertiport parcels varies. When  $A = 2$  acres, mean UAM travel times are lower than non-UAM, but as  $A > 2$ , mean travel times are at least 20 minutes higher for UAM than for all trips while driving.

Another limitation is the assumption of fixed 10-minute transfer times at the vertiports. In actuality, transfer times will likely vary based on the size and location of a vertiport. For example, it may take passengers longer to board and disembark at larger and more congested facilities. Because this could impact the comparative travel times between UAM and other modes of transportation, methodological advancements should attempt to more accurately estimate the time of the boarding, disembarkation, and transfer process based on the number of flights and passenger volumes at proposed vertiports [180].

The third limitation is the assumption of Marchetti's Constant, which may not always accurately reflect individual behavioral patterns [44]. Efforts to replace Marchetti's Constant involve UAM activity-based specifications to reflect more accurate individual behavioral patterns. This will allow for more accurate multimodal mode shares, as the model currently compares UAM vs. driving rather than all modes in the transportation network (e.g. public transit, biking, transportation networking companies, etc.). Several studies have already looked into this [85, 180, 233].

Another limitation is that the travel time analysis does not assume dynamic traffic assignment, where vehicles may change their routes in the middle of the trip or because of more recent congestion before embarking on the trip. This is currently an ongoing improvement in MANTA.

An important consideration for future work is to generalize the methodology for other metropolitan areas. The SF Bay Area has a strong central business district in San Francisco as well as a number of other large economic centers throughout the region (e.g., San Jose/Silicon Valley, Oakland, Tri-Valley, etc.). These economic centers contribute to a number of region-specific multidirectional travel patterns. As a result, the results of this study could be applied to metropolitan areas with both one large center (i.e., monocentric regions) and many multiple centers (i.e., polycentric regions) [214, 121]. The SF Bay Area also has a number of air traffic control (ATC) and infrastructure restrictions, which were not considered in this study but could be analyzed part of future research.

Additional work should also consider the cost of UAM which can impact mode choice [89, 85, 233]. Currently, the model may overestimate demand if UAM is more expensive than driving. Although more complex models are needed to compare UAM mode share alongside not only privately-owned vehicles but shared vehicle fleets (e.g., transportation network companies and automated taxis), the current model provides a methodological framework for commencing exploratory analysis.

## 4.4 Conclusion

While federal preemption currently governs most aspects of UAM regulation, establishing new UAM infrastructure is one area where local governments maintain a large amount of authority. While local governments cannot limit where UAM operates in the airspace, they can effectively limit where UAM aircraft take off and land through the zoning and siting

of vertiports. As such, land use planning, zoning, and regulation represent important tools that local governments can use to either support or hinder the growth of UAM.

Using a case study of the San Francisco Bay Area, we develop a vertiport design and network simulation tool to understand the travel time effects of UAM vertiport networks in the SF Bay Area under specific land use constraints. Given demand data at the zone level (further granularized to the node level) and extensive land use data at the parcel level, different UAM networks are developed with a varying number of vertiports and parcel constraints. These different configurations are developed using k-medians clustering and graph theoretic nearest neighbors. The vertiport locations are then integrated into the fast, parallelized GPU-based MANTA traffic simulator to understand travel times of UAM compared to driving. At  $k < 40$  vertiports at feasible locations, the same trips using UAM compared to driving have considerably higher travel times, but at  $k \geq 40$ , UAM's benefit in travel time emerges.  $k = 100$  vertiports is an inflection point in which the mean travel time for trips with both UAM and driving becomes lower than the mean travel time for only driving across all trips. In addition, as the area required for each parcel increases, travel times increase, implying that greater focus must be placed on efficient vertiport designs in smaller parcel areas. Although future research should incorporate additional modal connections (i.e., shared mobility and active transportation) to vertiports, intuitively a larger network of smaller vertiports could help reduce first- and last- mile travel distances and encourage non-vehicular connections, such as walking and micromobility.

While quantitative methods are an effective tool for assessing the technical viability of vertiport concepts, community engagement and local policy will be important to help understand potential community concerns, such as noise, aesthetics, traffic congestion, equity (i.e., displacement and gentrification), safety, and others. Community engagement will be an essential part of understanding and mitigating the adverse impacts of vertiports on vulnerable populations and communities. As part of the vertiport planning process, communities may consider embedding public participation and engagement at multiple stages. As part of long-range planning and visioning, communities may allow the public to discuss whether UAM is an appropriate strategy, and conduct exploratory scoping to understand potential UAM demand and the locations that could be compatible for vertiports. As the planning process progresses, communities may pivot toward more in-depth engagement processes that empower stakeholders to understand local concerns and use a variety of techniques to vet vertiport concepts and alternatives. If or when a developer decides to build a vertiport, additional community engagement should be sought and may be statutorily required by local, state/provincial, and federal laws.

Given the potential for UAM to grow over the next 10 to 20 years, communities may consider the potential impacts of UAM and planning considerations for multimodal integration and access/egress to vertiports. However, one challenge to establishing new UAM infrastructure is that UAM decentralizes the impacts of aviation that have typically been limited to airport facilities and their immediate surroundings. For this reason, UAM infrastructure and community integration will likely require close collaboration between aviation and community stakeholders that have traditionally worked in silos.

More research is needed to understand the potential impacts of vertiports on nearby communities, and identify design strategies and policies that could mitigate potential adverse impacts. Understanding the role of the built environment and its relationship to vertiport placement will be important to prioritizing sustainable transportation access and integration with nearby land uses. Communities may also need to develop policies for managing infrastructure access by multiple service providers, such as open, preferred, or exclusive access policies. The public sector may also need to consider the willingness to provide public support through mechanisms such as public-private partnerships. As the UAM marketplace continues to evolve, community planning, public outreach, and continued research will be important to understanding UAM's impacts on the transportation ecosystem and to balance community goals with commercial interests.



## Chapter 5

# Assessing the Value of UAM through Metropolitan-Scale Microsimulation

### 5.1 Introduction

Rapid advances in electric vertical takeoff and landing (eVTOL) aircraft, from higher battery energy densities to distributed electric propulsion, are driving the excitement in Urban Air Mobility (UAM). In parallel, the value proposition of UAM for a metropolitan area must be understood to inform future policy, infrastructure, and investment in the space. Even with enhancements of different driving modes, such as ridesharing and autonomous vehicles, increased surface-level congestion suggests there is a demand not met by the current transport modes; thus, UAM may fill the mid-to-long distance gap that the car is not efficiently or sustainably addressing [198, 185]. By supplementing inter-urban and intra-urban transportation demand with UAM, proponents anticipate that travel times will be greatly reduced, and long-run positive downstream effects may include lesser environmental degradation, less energy consumption, and fewer resources spent on infrastructure by leveraging free space [77, 22, 194].

However, to truly understand whether these potential benefits exist, we must simulate UAM travel in the transportation network. An initial evaluation of UAM compared directly against driving can provide insight. We compare against the automobile because it has dominated long-range metropolitan travel in the United States for decades, far beyond commuter rail and other public transit modes [11]. In addition, those with cars tend to have higher incomes, which will likely also be the initial user base for UAM before it becomes increasingly affordable over time and scale [90].

### 5.2 Background

UAM research for operations and infrastructure has insofar fallen in one of the following major categories: 1. Discrete choice modeling for UAM ridership, or 2. Vertiport network

design.

Surveys have been used to understand potential ridership and a traveler’s willingness to pay, finding that younger and higher income individuals are more likely to use UAM as a daily mode due to the greater time savings [90, 62, 92].

Vertiport network design has also been a popular subfield of UAM research, with a variety of approaches to find optimal vertiport locations, including weighted linear combinations [78, 179, 178, 181], spatial distributions of jobs and incomes [62], and k-means and k-medians clustering [77, 132, 203, 42, 237]. [181] used these designs to simulate UAM at a macroscopic level in various locales. These analyses found that UAM provides travel time savings for less than 1% of all trips and up to 15% if UAM adoption is high.

Many of these studies have presented stated preference discrete choice model results and/or simulation at a macroscopic level of granularity. However, no UAM research has, until now, has leveraged real-world data to determine demand, or high-fidelity, microscopic simulation to provide increased accuracy of vehicle behavior in congested regimes [181, 240]. In this paper, we determine UAM demand using an iterative data-driven approach under particular network designs, and subsequently microsimulate all trips at the regional scale, converging on a set of multi-modal UAM trips that benefit in travel time compared to driving.

### 5.3 Methodology

There are four major components of the methodology: (1) Data-driven multi-modal UAM demand generation, (2) Network design, (3) Driving and multi-modal UAM simulation, and (4) Convergence to a set of UAM trips that will benefit. We use the San Francisco (SF) Bay Area, one of the largest metropolitan regions in the United States, as the case study. The SF Bay Area is used for two reasons: (1) its economic and spatial growth patterns lead to both monocentrism and polycentrism, meaning that there was and continues to be one large economic center, San Francisco, but now there is also more dispersed economic activity across San Francisco, Oakland, San Jose, and peripheral cities such as Berkeley and Palo Alto [45, 103]. This makes for unique commuting patterns that could be useful in analyzing other cities of either monocentric or polycentric form. The second reason is functional, as data are available from the local metropolitan planning organization and transit agencies and many of the Silicon Valley mobility companies, such as Uber and Google.

#### Data-driven Multi-modal UAM Demand Generation

In this section, we describe a novel approach to determine the total addressable market of UAM by microsimulating real trips based on ridership and aerial flight constraints. This differs from the existing literature that uses stated preference surveys or static mode shares to determine multi-modal UAM demand.

The travel demand data are generated by the Bay Area Metropolitan Transportation Commission (MTC), the SF Bay Area’s metropolitan planning organization. MTC uses

its Travel Demand Model One to synthesize trips from census and household travel survey data [143], producing a full workday’s worth of trips on all modes. We filter the trips only to driving for use in MANTA. We filter further to the morning 5AM-12PM timeframe, producing approximately 3 million trips.

UAM demand is also based on aerial feasibility. A typical UAM quadrotor needs 2 miles to climb to its cruise altitude (2,000 ft), as specified in Fe<sup>3</sup> [235, 191], and shown in Figure 5.1.




Figure 5.1: A typical UAM flight trajectory, adapted from [191]

For two vertiports to be independent of each other, they must be separated by 4 miles. Thus, for any UAM quadrotor having a cruise segment, the origin and destination vertiports must be at least 6 miles apart for any feasible UAM mission. Section 5.3 shows that the average distance of a trip from 5AM-12PM is 7.08 miles, and the number of trips whose distances are 6 miles or greater is approximately 1.075M (37% of all trips). To understand the morning peak, we filter trips to the 8AM-9AM timeframe, which produces approximately 1.25M trips, with a 7.09 mile average distance. The number of trips that are 6 miles or greater in the 8AM-9AM slot is 481K (38% of all trips), as shown in Section 5.3. These proportions reflect that the distances of trips in the 8AM-9AM morning peak hour are representative of most trips throughout the full morning time period, suggesting that the 8AM-9AM hour for UAM traffic is not only relevant but also generalizable.

The Bay Area MTC breaks down the region into spatial units known as traffic analysis zones (TAZ). Each TAZ is a geographic unit that contains approximately 5,000 people, and thus its size can range from several blocks to an entire zip code, as shown in Figure 5.2. There are 1,454 TAZs in the SF Bay Area. Each trip has an origin and destination at the TAZ granularity. However, this is insufficient accuracy given the variance in areas of TAZs. Instead, we assign an origin and destination node to each trip by randomly sampling from a uniform distribution of the set of nodes contained in the respective TAZ. The centroid nodes are then calculated for each TAZ.

We now assume that every TAZ contains a vertiport, and that all vertiports can fly to all other vertiports, giving  $1,454^2$  possible vertiport pairs. We assume that every TAZ centroid is the location of that respective TAZ’s vertiport. Given the mission profile constraints from [235, 191], we filter the 8AM-9AM trips to those whose vertiport distances are  $\geq 6$  miles, producing approximately 400K trips. As a result, approximately 83% of trips have origin and destination vertiport distances that are  $\geq 6$  miles out of all the trips whose O to D distances are  $\geq 6$  miles. Section 5.3 shows that the number of trips between 8AM and 9AM exponentially declines as the inter-vertiport distance increases, preliminarily suggesting that




Figure 5.2: The travel analysis zones (TAZ) in the San Francisco Bay Area

UAM, which is expected to be valuable for longer trips, may address only a minority of trips rather than the majority of them.

## Network design

In this data-driven approach to determine the proper multi-modal UAM demand, the next step is the network design. Sample networks are created using k-medians clustering, partially adapted from [237]. However, in this case, the k-medians clustering algorithm uses the trips in the 8AM-9AM timeframe whose inter-vertiport distances are  $\geq 6$  miles, or approximately 400K trips. To understand the lower and upper bounds across the number of vertiports,  $k$ , we loop across the set of  $k = 10, 30, 50, 70, 100$  vertiports. The bulk of subsequent analysis will be with these vertiports and at transfer times of 2 minutes at each vertiport. Figure 5.4 shows the vertiport networks produced from the k-medians clustering algorithm at  $k = 10$ ,  $k = 30$ , and  $k = 100$ .

Once a network of vertiports is created for each  $k$ , the number of UAM-eligible trips for




Figure 5.3: In (a), the histogram of trips across distance from 5AM-12PM and from 8AM-9AM. The average distance and standard deviations are similar. In (b), the number of trips relative to the inter-vertiport distance for trips between 8AM and 9AM, showing a steep decline in trips as the distance increases.

that  $k$  is determined. This number will vary across  $k$ , as a trip is UAM-eligible if the origin and destination vertiports assigned from the  $k$ -medians clustering algorithm are different. For instance, at a low  $k = 10$ , there will be more trips whose origin and destination vertiports are the same because there are so few, highly spaced out vertiports available. However, at  $k = 100$ , there will be more UAM-eligible trips because there is a higher likelihood that a trip’s origin and destination vertiport are different. Figure 5.5 shows that at  $k = 10$ , there are 235K UAM-eligible trips out of 400K. At  $k = 30$ , there are 341K UAM-eligible trips out of 400K. At  $k = 100$ , there are nearly 387K UAM-eligible trips out of 400K. To satisfy the aerial mission profile, we see that all UAM-eligible trips at  $k = 10$  have O and D vertiport distances that are  $\geq 6$  miles. At  $k = 30$ , there are 333K trips out of 341K, or 98% of trips. At  $k = 100$ , there are 345K trips out of 387K trips, or 89% out of the original UAM-eligible trips. The trips that do not satisfy this constraint are reassigned as driving trips.

## Driving and Multi-Modal UAM Simulation

In order to simulate both driving and multi-modal UAM trips, we use Microsimulation Analysis for Network Traffic Assignment (MANTA), the fast, GPU-based traffic microsimulation platform, developed by this author [240]. Existing simulators have typically revealed a tradeoff among accuracy, computational speed, geographic scale of simulation, and cost [240]. However, MANTA exhibits performance benefits in all of these areas, enabling it to be used for agile scenario planning, such as with different UAM network designs and iterative convergence techniques. In addition, MANTA’s ability to track the microscopic movements of

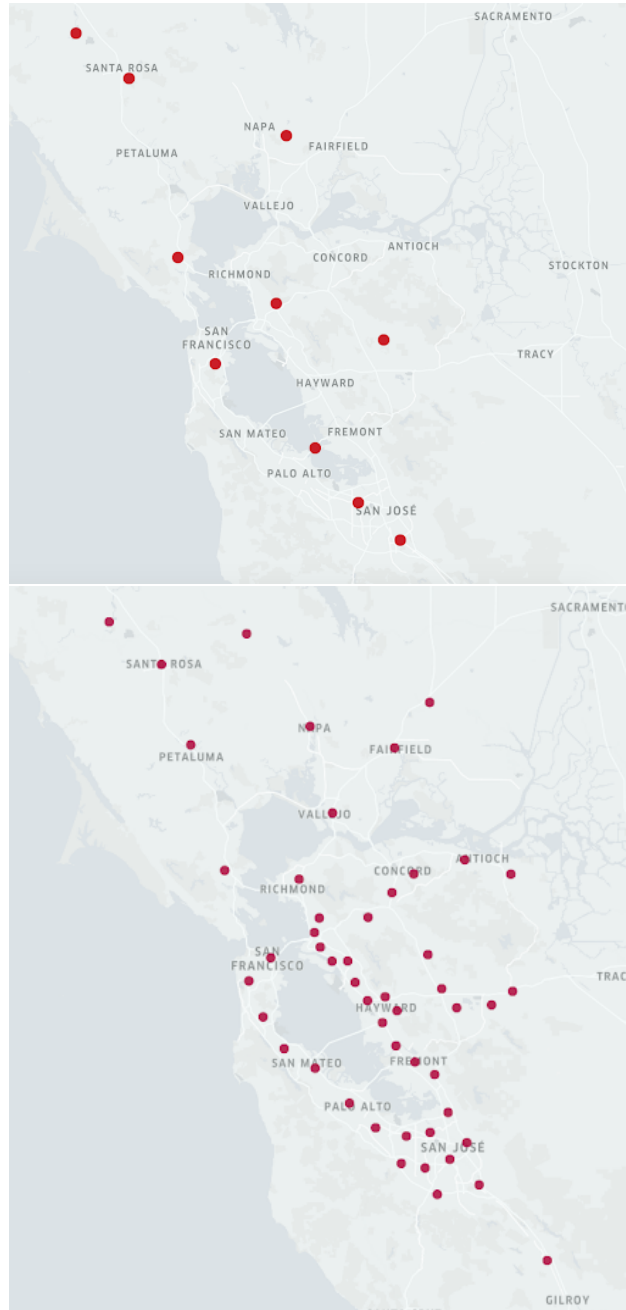


Figure 5.4: Network designs determined from k-medians clustering at  $k = 10$  and  $k = 100$




Figure 5.5: Distances

vehicles on the street network, namely lane changes, acceleration, and braking with respect to other vehicles, enables a granularity and accuracy vital to travel time computations [240].

There are several steps in the simulation process. The first step is to simulate all 3M trips in the 5AM-12PM timeframe as driving trips. This produces driving times for every trip, which are later used as comparison in the travel time benefit analysis. This total number of trips will be generalized as  $D$  in the rest of the paper.

For each  $k$ , a set of trips is determined as UAM-eligible, highlighted in the previous section. Each of these trips is then broken down as a multi-modal UAM trip, which contains three legs as well as transfer times at the vertiports: 1. Origin to its nearest origin vertiport (access), 2. Origin vertiport to destination vertiport (aerial flight), and 3. Destination vertiport to the destination (egress), as shown in Figure 5.6.




Figure 5.6: Every multi-modal UAM trip contains an access leg, transfer at the origin vertiport, aerial leg, transfer at the destination vertiport, and an egress leg.

The total travel time of each multi-modal UAM trip is shown in Eq. (5.1)

$$T_{trip} = T_{o,o_v} + T_{o_v} + T_{o_v,d_v} + T_{d_v} + T_d \quad (5.1)$$

where  $T_{trip}$  is the total travel time of the multi-modal UAM trip,  $T_{o,o_v}$  is the travel time via automobile from the origin node to the origin vertiport,  $T_{o_v,d_v}$  is the aerial travel time between vertiports,  $T_d$  is the travel time via automobile from the destination vertiport to the destination node, and  $T_{o_v} = T_{d_v}$  is the transfer time at each vertiport, which assumes a buffer time for security, loading, and embarking (and vice versa).

For each  $k$ , this set of UAM trips,  $U$ , produces  $2U$  car trips due to the presence of both access and egress driving legs for each multi-modal UAM trip. This is then combined with the rest of the demand  $D - U$ , which remain driving trips, producing a total of  $D + U$  trips on the road. All of these trips are then simulated and processed in MANTA, and the total travel times for the  $U$  trips are compared to the driving times for the same trips, as will be shown later in the paper.

## Convergence

While at each  $k$  the UAM-eligible trips are the ones offloaded to multi-modal UAM, this does not take into account whether the user would *choose* multi-modal UAM over driving at all. If this occurs, then simulating all the travelers based on that number of trips being offloaded to UAM - while the others in the demand remain driving trips - will produce results in which some travelers will wish to change their mode.

User equilibrium is a common problem in transportation simulation, in which one traveler's mode choice can affect the travel times of other travelers, who may then change their mode choice due to the change in travel times [198]. For this reason, a user equilibrium (UE), or Wardrop's first principle, must be satisfied in which no traveler's travel time on their current route can decrease by changing to a different route or mode. This problem can be addressed by using an iterative convergence algorithm, as shown in Equation (5.3).




Figure 5.7: Iterative convergence across multiple simulations narrows down the number of trips that benefit from taking multi-modal UAM

Assuming travel time is the only parameter in our mode choice model, the traveler would choose multi-modal UAM over driving if and only if its travel time with multi-modal UAM is lower than that of driving. This convergence process is best explained with an example. For instance, if there are 3M total trips, and 10% is offloaded to UAM, then 300K trips will take



UAM and 2.7M trips will drive. Once this is simulated, we will notice that not all 300K trips may benefit from using UAM (e.g., their travel times are, in fact, longer with multi-modal UAM than driving). In this case, if 10% of the 300K trips did not benefit from UAM, then there are 30K trips that did not benefit and 270K trips that did benefit. However, this benefit was based on when there were 300K trips offloaded to UAM on the network while 2.7M trips were still driving. The number of trips that benefit may change if we are to simulate again with only the number of trips that benefited from UAM with the rest assigned to drive. In other words, we must now re-simulate with 270K trips that take multi-modal UAM while 2.73M trips drive. As this iterative process continues, the delta in the number of trips that benefit between each iteration will eventually decrease below a certain threshold, signifying that the system converged to a user equilibrium of all travelers choosing the mode that best minimizes their travel times. The process is formalized in Eq. (5.2) and (5.3).

$$\beta_i = T_{d,i} - T_{u,i} > 0 \quad (5.2)$$

where  $\beta$  is the travel time benefit for trip  $i$  when taking UAM compared to driving,  $T_{d,i}$  is the travel time for trip  $i$  while driving, and  $T_{u,i}$  is the travel time for trip  $i$  while driving. Let  $N_t$  be the total number of trips in  $\{\beta_1, \beta_2, \dots, \beta_i > 0\}$  for iteration  $t$ . Eq. (5.3) then specifies the iterative convergence criterion in which all travelers are taking the mode that produces travel time benefit compared to the other mode.

$$N_{t+1} - N_t \leq \eta \quad (5.3)$$

where  $t$  is the iteration, and  $\eta$  is the threshold at which the system stabilizes, representing the delta between the number of benefited trips between two successive iterations.

Generalizing, let  $D$  be the total number of trips and  $N_1$  the number of trips initially assigned to multi-modal UAM in iteration 1. Thus,  $D - N_1$  trips are assigned as driving trips and  $N_1$  trips are assigned as multi-modal UAM trips. Then, once all trips are simulated in iteration 1, only  $N_2$  trips actually benefit from taking multi-modal UAM. Then, in iteration 2,  $D - N_2$  trips are assigned as driving trips and  $N_2$  trips are assigned as multi-modal UAM trips. From the results of iteration 2, it may show that  $N_3$  trips benefited from taking multi-modal UAM. Thus, in iteration 3,  $D - N_3$  trips are assigned as driving trips and  $N_3$  trips are assigned as multi-modal UAM trips. This process continues and is reflected in Eq. (5.3). Ultimately, we expect that  $N_1 \geq N_2 \geq N_3 \geq \dots N_T$ , where  $T$  is the total number of iterations required to converge. The differences between unconverged and converged trips are shown in the next section.

## 5.4 Results

For each  $k$ , the entire simulation process, including convergence, is carried out. Due to the speed of MANTA, which is parallelized, each iteration takes under 6 minutes [240]. Such speed and performance allows for fast scenario planning and analysis across different parameters.

## Transfer Times

As previously mentioned, an important consideration is the transfer time at the origin and destination vertiports for each multi-modal UAM trip. From Figure 5.8, we can see that the number of benefited trips across  $k$  is highly sensitive to the transfer time. At a transfer time of 2 minutes at each vertiport, approximately 10% of the UAM-eligible trips benefit (post-convergence) at  $k = 10$ , 35% at  $k = 30$ , and 65% at  $k = 100$ . The largest jump in share of trips occurs between 10 and 30 vertiports, and the rate of improvement slows from 30 to 100 vertiports. At transfer times of 5 and 10 minutes, the number of benefited trips decreases successively by approximately 6K trips at  $k = 10$ , 30K trips at  $k = 30$ , and 50K trips at  $k = 100$ . For all  $k$ , the share of trips decreases at a higher rate when the transfer time changes from 5 minutes to 10 minutes compared to when it changes from 2 minutes to 5 minutes. This suggests that the margin of benefit for many multi-modal UAM trips, regardless of  $k$ , is likely around 10 minutes (the delta between two 10 minute transfer times vs. two 5 minute transfer times). In Figure 5.8, the number of benefited trips across  $k$  follows a logarithmic pattern across all transfer times.




Figure 5.8: The number of trips that benefit from multi-modal UAM

## Convergence

For the following convergence analysis, the transfer time at each vertiport is assumed to be 2 minutes. We select  $\eta = 10$  trips, such that there is no more than a 10 trip delta between the number of benefited trips in the previous iteration compared to the current iteration. Across all  $k$ , a range of 2 to 5 iterations were required. Figure 5.9 shows that at

$k = 10$  and  $k = 30$ , the number of benefited trips after convergence is almost 200K trips fewer than before convergence. At  $k = 100$ , the number of converged trips differs from the number of unconverged trips by approximately 130K. With convergence, the trips that benefit have shorter multi-modal UAM travel times than those in the unconverged scenario, which produces higher deltas between the multi-modal UAM travel time vs. driving for the same trip.

From Section 5.4, the average distance of the trips that benefited after convergence shows that higher distance trips, around 54 km, benefit more at the lower  $k = 10$ , and at  $k \geq 30$ , medium-distance trips below 40 km begin to benefit as well.

In Section 5.4, the travel time delta,  $\delta$ , represents the difference in travel time between driving and multi-modal UAM for the same trip. When  $\delta < 0$ , the driving time is lower than the multi-modal UAM time; when  $\delta > 0$ , the multi-modal time is lower than the driving time, and it is considered a trip that benefits. Across  $k$ , the average travel time delta stays relatively stable around 20 to 23 minutes, which confirms that both multi-modal UAM and driving times are still sensitive to the distance of the trips themselves.

## Regional Analysis

For any new emergent mobility mode to gain traction for investment and infrastructure provision, its impact on the existing transportation network must be evaluated. Specifically, metropolitan planning organizations must understand the impacts to road congestion if a new mode is introduced. The prevailing thought among proponents of UAM is that UAM will be able to offload a sufficient number of trips such that the regional congestion on the ground is noticeably improved. Indeed, this is generally reflected in the results, with one caveat. The results in Figure 5.10 shows that changes in road congestion are highly dependent on the vertiport network design, and particularly on  $k$ . As shown in Figure 5.10, and assuming a transfer time of 2 minutes, at  $k = 10$  and  $k = 100$ , the average travel time of all trips on the road is lower than driving by approximately 1.5 to 5 minutes, respectively, which is significant at an aggregate scale.

With the decreases in overall travel times and congestion, we must also address whether incorporating a new mode is fair for those who never use it. In Figure 5.11, we find that those who never take UAM very marginally benefit, from on average about 1/10th of a minute at  $k = 10$  to about 1 minute at  $k = 100$ . There is also a larger increase in benefit from 10 vertiports to 30 vertiports, suggesting that there may be a sweet spot at a minimum of 30 vertiports where all travelers, both multi-modal UAM and non-UAM users, benefit noticeably.

## 5.5 Limitations and Future Research

An important finding shows that travel times are extremely sensitive to exact network designs. When performing k-medians clustering, the algorithm chooses random points within

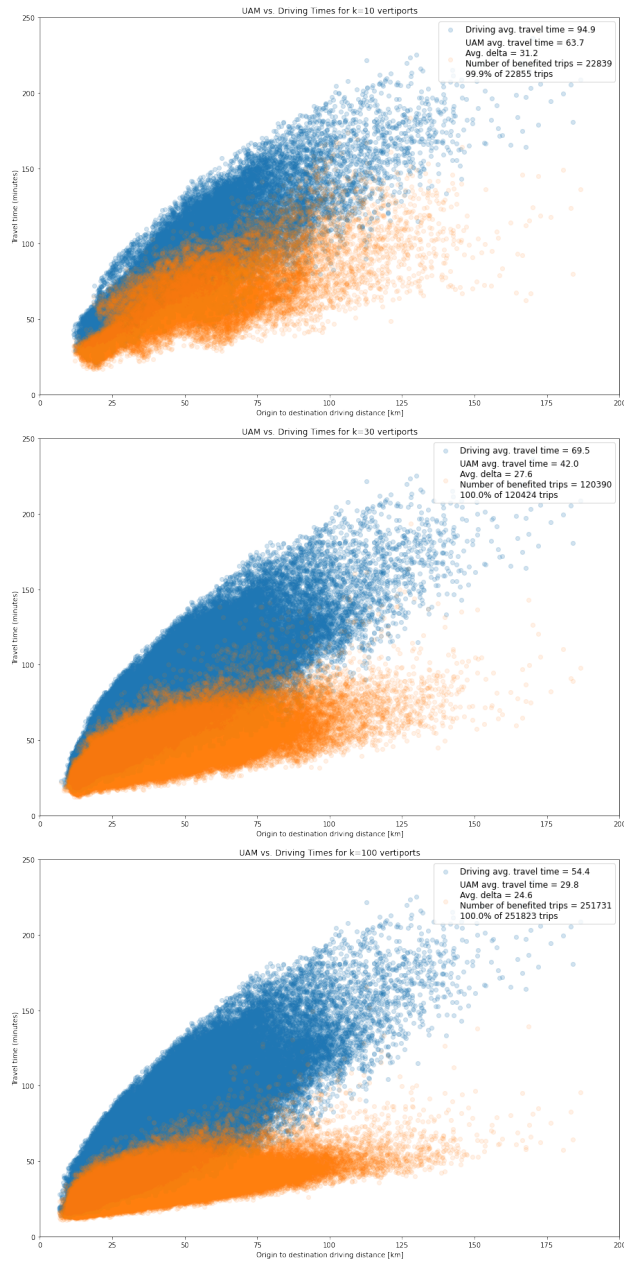


Figure 5.9: The number of benefited trips across  $k = 10, 30, 100$  after convergence




Figure 5.10: At  $k \geq 10$ , the average travel time on the road is lower than driving those same trips, showing that there is benefit to having UAM in the transportation network once there are at least 10 vertiports, with significant benefits occurring at 30 vertiports or more.

the street network that optimizes the objective. However, it does not consider whether these nodes have intersecting edges that are small tertiary roads that could overflow with excess congestion due to so many trips traveling to and from these few vertiports. An improvement would be to actively select locations that are suitable for a vertiport as the vertiport candidates, considering land use parameters, ground congestion, and vehicle capacity of nearby edges [189].

These results also reflect that vertiport capacities need to be incorporated into network design and analysis, as there is severe demand imbalance across vertiports. As shown in Figure 5.12 at  $k = 10$  vertiports, there are several vertiports servicing hundreds to even one thousand trips more than other vertiports. This imbalance also exists at higher  $k$ .




Figure 5.11: Trips that do not take UAM also benefit, albeit slightly, when UAM is an available mode in the transportation network.




Figure 5.12: At  $k = 10$ , there is an imbalance among vertiports in how many trips are served

In addition to more optimal network design, further research should incorporate more realistic constraints of the ground-air interface at the vertiport, considering both vertiport and aircraft capacities, which will likely have a significant effect on total multi-modal UAM travel times.

## 5.6 Conclusion

This study provides a novel regional-scale microsimulation analysis of UAM's benefit in a metropolitan area, using the San Francisco Bay Area as a case study. We first present a data-driven approach to determining UAM demand, and then subsequently simulate multi-modal UAM along with driving trips across the entire metropolitan area using the fast, highly parallelized regional-scale microsimulator MANTA. We find that the short distance trips are gained by expanding the network to a higher  $k$ , and the long distance trips benefit across all network designs. These benefits, though, are sensitive to transfer times at the vertiports, suggesting that vertiport scheduling and capacity may become bottlenecks that limit the value of UAM. We also find that at  $k \geq 10$  vertiports, the total regional congestion with UAM as a mode is lower than if there was no UAM and all trips were driving trips, with noticeable improvement to non-UAM travelers at  $k \geq 30$  vertiports. Even considering suboptimal network design, this shows that UAM can have significant benefit for a metropolitan region.

# Chapter 6

## SimUAM: A Regional-Scale Multimodal UAM Toolchain

### 6.1 Introduction

When a new mode is introduced into a metropolitan transportation network, its impact on individual travel behavior is often not readily understood [97]. To proponents, UAM directly tackles one of the known parameters of mode choice, travel time, a theoretically and empirically known ‘bad’, or a factor that increases disutility [199, 89]. Assuming UAM operations will use an on-demand model that includes aircraft flying along high-volume travel corridors between a network of vertiports, analyses have attempted to estimate the potential impacts of UAM passenger services on travel behavior, using predetermined vertiport locations and simulating different modes of travel in different metropolitan areas [180, 216, 42, 237, 132]. However, these existing models assume static parameters in either ground traffic, vertiport transfer times, aerial flight, or a combination of the three. For UAM, since its purported value is in travel time savings, simulating the margins at a granular level, incorporating exact dynamics and scheduling in each leg of the trip, is of utmost importance.

This paper proposes SimUAM, a holistic multi-modal toolchain that integrates ground traffic microsimulation, ground-air interfacing at the vertiport, and aerial microsimulation, producing exact travel times for both multi-modal UAM trips as well as other driving trips across the region. Specifically, the MANTA component of SimUAM models road congestion patterns of all driving trips, including the access and egress legs of the multi-modal UAM trips. VertiSim models the exact ground-air interface at the vertiport to move the traveler from the ground to a specific aircraft for the multi-modal UAM trips. The  $FE^3$  component models the exact flight trajectories of these aircraft, incorporating conflict resolution and real UAV mission profiles in the sky. For each multi-modal UAM trip, the *interaction* among the three legs is exceedingly important in realizing the true value of UAM in a regional transportation network. As a result, SimUAM is the first holistic toolchain, to the best of these authors’ knowledge, that models all these interactions together dynamically and






Figure 6.1: Simulation components for an individual multi-modal UAM trip

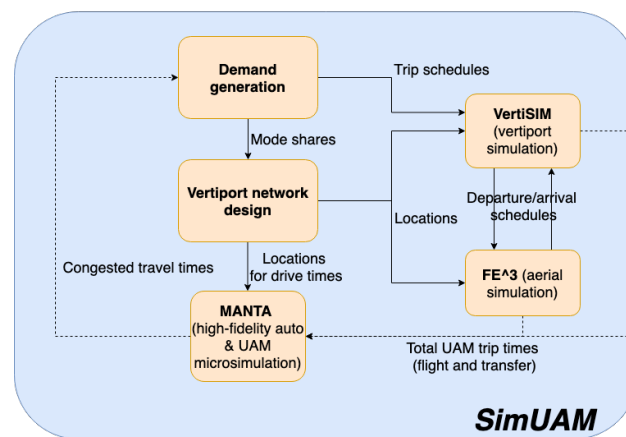


Figure 6.2: The full SimUAM model pipeline

granularly, as shown in Figure 6.1.

This paper is organized into the following sections. Section 6.2 details MANTA, the GPU-based regional-scale road traffic microsimulator. Section 6.3 explains VertiSim, the granular, discrete-event vertiport and pedestrian simulator. Section 6.4 discusses  $FE^3$ , a NASA-developed, GPU-based high-fidelity regional-scale microsimulator for air traffic. These three components comprise the SimUAM architecture, shown in Figure 6.2. Section 6.6 describes the performance metrics measured by SimUAM and experiments showing SimUAM’s benefit. Finally, Section 6.7 and Section 6.8 highlight the limitations, future work, and conclusions of this research.

## 6.2 MANTA: Microsimulation Analysis for Network Traffic Assignment

The SimUAM framework relies on accurate ground traffic microsimulation in order to produce congested travel times for all the travel demand in the investigated region. In order

to do so, we leverage Microsimulation Analysis for Network Traffic Assignment (MANTA), an ultra-fast, highly-parallelized GPU-based microsimulation platform, developed by this author [240]. Existing simulators have typically revealed a tradeoff among accuracy, computational speed, and geographic scale of simulation [240, 15, 46, 165, 175]. However, MANTA exhibits performance benefits in all three of these areas, enabling it to be used for agile scenario planning, particularly with an emergent mode such as UAM, whose deployment is still in its inchoate stages. MANTA’s ability to track the microscopic movements of vehicles on the street network, namely lane changes, acceleration, and braking with respect to other vehicles, enables a granularity that remains vital in travel time computations [240].

## Inputs

MANTA takes two major inputs: (1) Geographic network, comprised of edges and nodes in a metropolitan geography’s street network, and (2) Origin-destination (OD) demand showing traveler movements in the selected geography.

The geographic network, as described in [240], is derived from OpenStreetMap. We use the network library OSMnx to determine the nodes that exist within the polygonal hull of a metropolitan region, determined from widely available shapefiles for the region of investigation [34]. The San Francisco (SF) Bay Area is shown in Figure 6.3 as an example output of this process, and is also used as the area of investigation in Section 6.6 of this paper. The SF Bay Area is chosen as the case study due to wide travel data availability and massive geographic scale that reflects useful bidirectional commuting patterns.

The OD demand is generated from the local metropolitan planning organization (MPO), which typically calibrates a travel demand model based on estimates from census and household travel data [143]. For the SF Bay Area, the data are derived from the Bay Area Metropolitan Transportation Commission (MTC)’s Travel Demand One model.

Once each trip is assigned a mode, these trips are filtered to only driving trips, which include driving alone, shared trips, and driving to transit. In the SF Bay Area, this totals 25M trips for the full day. This work focuses on the morning commute, from 5am-12pm, and is filtered down to approximately 3M trips. Importantly, these trips may include multiple passengers in the car, but these extra passengers are not considered in the number of trips.

The origin and destination of each trip is at a traffic analysis zone (TAZ) granularity, which is a population density-based geographic unit used by MPOs across the nation, with each TAZ area typically larger than a block but smaller than a zip code. This size is too coarse for MANTA, which operates at the node level. As a result, once the origin and destination TAZs are known, we randomly assign nodes to the origin and destination within their respective TAZs, sampled from a uniform distribution, producing each trip as one from an origin node to a destination node.




Figure 6.3: The San Francisco Bay Area region’s edges and nodes, defined by the polygonal hull of its nine counties

## Functionality

In MANTA, there are two major components: (1) routing and (2) microsimulation. The initial routing algorithm developed in MANTA used a parallelized Dijkstra priority queue implementation, which greedily selects the closest vertex that has not yet been processed. This single source shortest path (SSSP) implementation could compute 3M routes, the number of driving trips from 5 AM to 12 PM in the SF Bay Area, in approximately 1 hour.

However, an update was made to integrate an ultra-fast, parallelized Open Source Routing Machine framework using Pandana, based on the contraction hierarchy scheme [84]. The time it takes to compute the same 3M routes in the SF Bay Area is now under one minute.

MANTA is a time-based microsimulator that accounts for parallel changes that occur due to the dependence of cars’ behaviors on one another. The vehicular movement on an edge is dictated by conventional car following, lane changing, and gap acceptance algorithms [240, 210]. The well-known Intelligent Driver Model (IDM), as shown in 6.1, is used to control the vehicle dynamics through the network [213].

$$\dot{v} = a \left( 1 - \left( \frac{v}{v_o} \right)^\delta - \left( \frac{s_o + Tv + \frac{v\Delta v}{2\sqrt{ab}}}{s} \right)^2 \right) \quad (6.1)$$

where  $\dot{v}$  is the current acceleration of the vehicle,  $a$  is the maximum possible acceleration of the vehicle,  $v$  is the current speed of the vehicle,  $v_o$  is the speed limit of the edge,  $\delta$  is

the acceleration exponent,  $s$  is the gap between the vehicle and the leading vehicle,  $s_0$  is the minimum spacing allowed between vehicles when they are at a standstill,  $T$  is the desired time headway, and  $b$  is the braking deceleration of the vehicle [240, 227]. The exact position of each vehicle at the current timestep is computed using this calculated acceleration value  $\dot{v}$ .

In addition to car following, vehicles also follow lane changing rules, based on whether the vehicle is making a mandatory or discretionary lane change. Mandatory lane changes occur when the vehicle must take an exit from the edge, while discretionary lane changes occur during overtaking a slower vehicle [240]. The lane changing model states that the vehicle has an exponential probability of switching from a discretionary lane change to a mandatory lane change as it proceeds through the edge, as shown in 6.2.

$$m_i = \begin{cases} e^{-(x_i-x_0)^2} & x_i \geq x_0 \\ 1 & x_i \leq x_0 \end{cases} \quad (6.2)$$

where  $m_i$  is the probability of a mandatory lane change for vehicle  $i$ ,  $x_i$  is the distance of vehicle  $i$  to an exit or intersection, and  $x_0$  is the distance of a critical location to the exit or intersection [111, 236].

Once a vehicle decides to change lanes, the maneuver is performed if the lead and lag gaps of the cars in the lane to which it is changing are acceptable. The critical lead or lag gap for a successful lane change is defined as the minimum distance to the following or lagging vehicle that allows for a lane change, as shown in 6.3.

$$g_{lead} = \max(g_a, g_a + \alpha_{a1}v_i + \alpha_{a2}(v_i - v_a)) + \epsilon_a \quad (6.3)$$

$$g_{lag} = \max(g_b, g_b + \alpha_{b1}v_i + \alpha_{b2}(v_i - v_b)) + \epsilon_b \quad (6.4)$$

where  $g_{lead}$  is the critical lead gap for a lane change,  $g_{lag}$  is the critical lag gap for a lane change,  $g_a$  is the desired lead gap for a lane change,  $g_b$  is the desired lag gap for a lane change,  $\alpha$  is a system parameter (typically [0.05,0.40]) that controls the gap based on speed,  $v_i$  is the speed of the vehicle,  $v_a$  is the speed of the lead vehicle,  $v_b$  is the speed of the lag vehicle, and  $\epsilon_a$  and  $\epsilon_b$  are the random components [240].

MANTA is a parallelized, time-based simulator, rather than an event-based simulator, and thus accounts for changes that occur for all vehicles in every timestep. A time-based approach makes parallelization easier to implement, as synchronization protocols for event-driven simulators is challenging. In addition, no overhead exists with time-based simulators, which contrasts with event-based simulators, where millions to billions of events must be constantly generated. Without this overhead, this allows for the ability to model more granular movements in edges, not just events such as a vehicle entering or exiting an edge, which often constitutes an event. The MANTA architecture is shown in Fig 6.4.

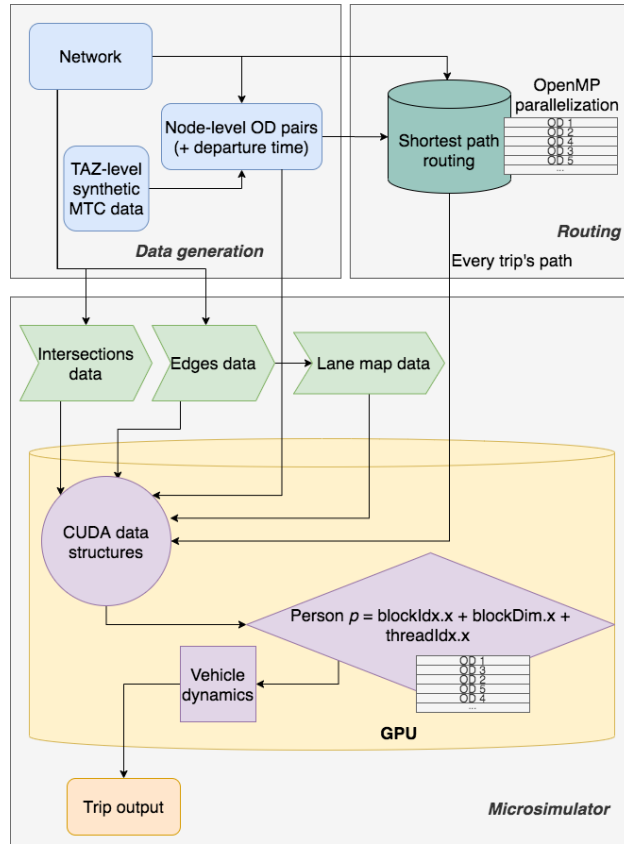


Figure 6.4: The MANTA architecture

## Calibration and Validation

MANTA has been calibrated against Uber Movement data on a per-edge basis for the San Francisco Bay Area. First, 17% of edges in the San Francisco Bay Area network are matched with the Uber Movement data [240]. Then, a mini-batch gradient descent algorithm is used to iteratively converge to the optimal  $a$ ,  $b$ ,  $T$ , and  $s_0$  parameters of the Intelligent Driver Model, shown in 6.1. In addition, each driver is outfitted with a different driver profile per edge, based on a Gaussian distribution, with the speed limit of the edge as the mean and the standard deviation of the Uber speeds on the edge as the standard deviation. This produces closely matched edge speeds for the 17% of edges, as shown in Figure 6.5 [240]. Average travel times and speeds are then closely validated with California Household Travel Survey (CHTS) data as well as Uber Movement data [240]. As an example, Figure 6.5 shows the mean edge speeds across all edges whose speed limits are 45 mph, for both MANTA and Uber Movement. The average speeds differ by 3 mph and the standard deviations by below 1 mph, suggesting successful calibration.




Figure 6.5: Kernel density estimator plot comparing MANTA microsim edge speeds vs. Uber edge speeds on edges whose speed limits are 45 mph

From a performance perspective, MANTA can run one full simulation with 3M trips, including routing, in 4.6 minutes on a machine with Intel i9-7940X CPU, 3.10 GHz clock frequency and 28 cores. The GPU that runs the microsimulation component is an NVIDIA GP104 (GeForce GTX 1080) with a 33 MHz clock.

MANTA provides the simulation backbone to the entire SimUAM platform, as all trips, both driving and multi-modal UAM, must originate and end through MANTA. More information can be found in [240]. However, as shown in the next section, multi-modal UAM trips must be modeled by multiple interacting components, one of which is the ground-air interface modeled by VertiSim.

### 6.3 VertiSim: Vertiport Ground-Air Interface Modeling

Aircraft need infrastructure to land, taxi-in, park, load and unload passenger/cargo, charge, repair, taxi-out, and take-off. Vertiports have to provide this infrastructure to handle passengers and aircraft within limited land use. Thus, they have limited throughput capacity. This limitation makes vertiports the bottlenecks of the UAM system [219].

Vertiports are the transfer centers for passengers to change from one mode to another. Therefore, transfers at the origin and destination vertiports must be simulated in order to accurately calculate their transfer time. Additionally, aircraft turnaround time and taxiway conflicts must be simulated for accurate throughput calculation.

In order to understand passenger waiting times and aircraft throughput, VertiSim has been developed by the second author. VertiSim is an agent-based discrete-event simulator capable of simulating millions of passengers and aircraft.

## Inputs

VertiSim takes as input vertiport locations and surface layouts (topologies). These are comprised of TLOF pads, parking spaces, number of charging stations, number of security lines/turnstiles at the vertiport entrance, distributions for stochastic walking times for passengers, passenger arrival schedule with passengers' destination vertiport, aircraft arrival schedule. It also takes several inputs for vertiport surface management strategies.

## Functionality

The three major components of VertiSim are the (1) vertiport layout designer, (2) flight generator, and (3) vertiport manager. The responsibility of the vertiport layout designer is to create the node-link model of the input file. The flight generator then generates flights, as the number of passengers accumulate to fly to the same direction, and assigns passengers to available aircraft. If there is no available aircraft, the generator puts them in the departure queue. It also generates flights for the passengers who waited more than  $T_{max}$ , the maximum waiting time threshold, specified intuitively as 10 minutes, without waiting for the aircraft to be full. When there is a flight queue for the departing passengers in the waiting room, they wait until their turn. The flight generator can assign new passengers to these underbooked flights if passengers who are flying to the same direction arrive to the waiting room before passengers leave the waiting room. The vertiport manager oversees the usage of the resources of the vertiport for the agents. That is, it assigns a parking space and TLOF pad for the aircraft and creates the taxi route for the aircraft. Ultimately, VertiSim is able to simulate the departing passenger simulation, arriving passenger simulation and aircraft simulation at the vertiport surface.

VertiSim is capable of simulating a variety of vertiport topologies. Every vertiport in the given vertiport network can have a unique topology. Departing passenger flow in VertiSim is as follows: (1) car to vertiport entrance, (2) security check/turnstile pass, (3) security check/turnstile to waiting room, (4) waiting room to boarding gate, (5) boarding gate to aircraft, (6) taxi-out, (7) take-off. Arriving passengers then disembark, go to the exit gate at the vertiport surface, exit the vertiport at the ground level, and enter a car. From the perspective of the aircraft, arriving aircraft request an available parking space. If there are no available parking spaces, they loiter until one is reserved. They then land at the TLOF pad assigned by the vertiport manager, taxi-in, unload the passengers, and finally charge. These aircraft then become available for the departing passengers. When a parked aircraft is assigned to a flight, they load passengers, reserve an available TLOF pad, taxi-out, and take-off.

VertiSim outputs aircraft departure schedules, aircraft turnaround time statistics, passenger transfer time statistics, and vertiport efficiency metrics. The VertiSim architecture is shown in Figure 6.6.

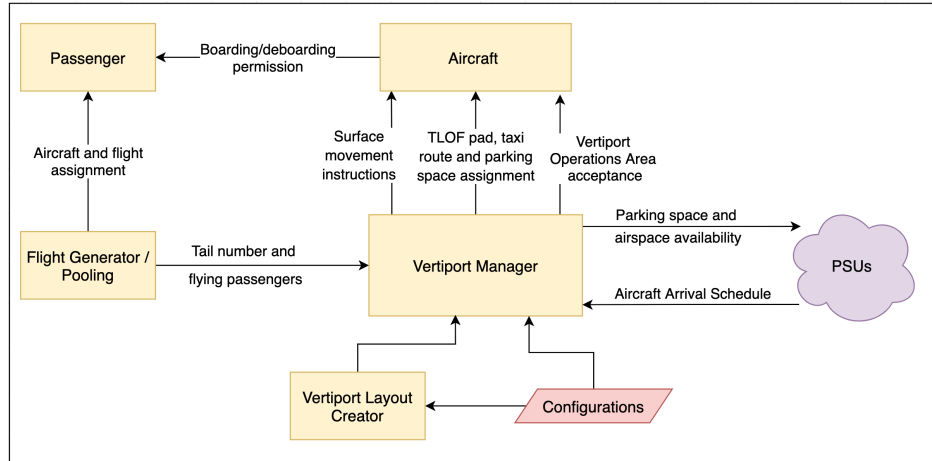


Figure 6.6: The ground/air interface at the vertiport

### Steady-state data collection

The first passengers have priority to fly when the minimum number of passengers required to fly to the same direction has been met, or if they have waited more than  $T_{max}$ . Therefore, the first agents wait less than the other agents. In order to collect accurate data, we remove the data from the transient-state. Although there is no general guideline to identify transient-state from steady-state, several methods exist. The first option is to choose a large enough simulation time to eliminate the effects of the transient state. The second option is to calculate across independent simulation runs the beginning time of the transient-state [37]. We have derived a formulation to detect the beginning of the steady-state region based on our observations.

$$N_A = N_S * N_{TLOF} * 2 \quad (6.5)$$

where  $N_A$  represents number of pseudo-agents need to be added on top of the original demand of each vertiport,  $N_S$  is aircraft passenger capacity, and  $N_{TLOF}$  represents number of TLOF pads at the vertiport. For instance, on a 4 seat aircraft and 4 TLOF pad scenario, VertiSim creates 32 passengers before the actual demand arrives. Pseudo-agents are randomly sampled from the original demand file and removed at the end of the simulation.

### Calibration and Validation

The parameters of VertiSim such as average human walk speed, aircraft ground speed are determined from the literature and by subject-matter experts [129, 244]. Comparing the location-time graphs of the passenger and aircraft agents to these studies confirm the proper parameters.



## 6.4 $FE^3$ : High-fidelity Aerial Simulation

In order to model UAM accurately, we must also simulate aerial component of the trip. Leveraging the  $FE^3$  GPU-based microsimulation platform, developed at NASA, a mission profile is created for every UAM trip, with a departure time, ascent and descent characteristics, cruise altitude, and velocity/directional changes depending on wind patterns and collision avoidance [235]. The  $FE^3$  platform is able to model these in-flight dynamics and any subsequent flight trajectory shifts that contribute to changes in total flight time, making it ideal for simulating UAM in congested airspace [235].  $FE^3$  is a dynamic, high-fidelity microsimulator leveraging a parallelized GPU architecture, similar to MANTA. Due to the similarities in fidelity, the loose coupling is both seamless and interpretable.

### Inputs

$FE^3$  takes several structures as input. As described earlier, MANTA outputs the arrival times of the multi-modal UAM trips to their respective origin vertiports, and VertiSim consolidates these passengers into aircraft based on pedestrian walking times, aircraft arrival times, and vertiport capacities. Once these passengers are assigned aircraft and then board, the aircraft must go through the takeoff and landing process before it departs. As a result,  $FE^3$  takes as input the aircraft departure times output from VertiSim. The passenger IDs are abstracted away from  $FE^3$ , which proceeds based on the assumption that all aircraft contain at least 1 passenger. In addition to the aircraft departure times,  $FE^3$  defines simulation related parameters, such as the number of Monte Carlo simulations, temperature and humidity, and the type of avoidance algorithm. It also has the ability to overwrite certain vehicle or device model parameters, as shown in Figure 6.7.

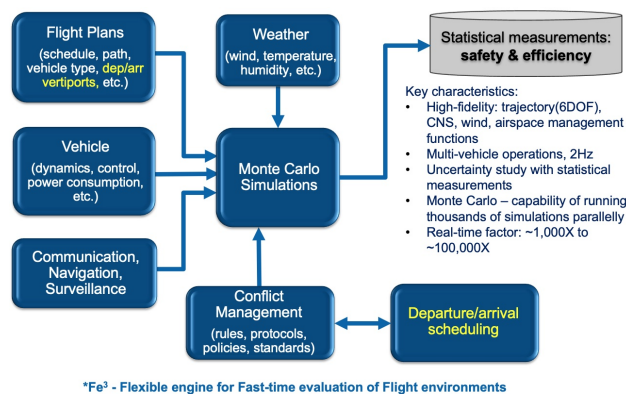


Figure 6.7:  $FE^3$  aerial microsimulator architecture

## Functionality

FE<sup>3</sup> is highly-parallelized using the CUDA programming language on graphics processing units (GPUs). It is deployed on the Amazon Web Service (AWS) cloud for scalability, such that the number of GPU instances can be dynamically deployed based on simulation needs. Within the simulator, each aircraft's flight in FE<sup>3</sup> is governed by several models, specifically trajectory, energy consumption, collision avoidance, vehicle communication and sensors, and wind models [235].

The trajectory model consists of simulating the vehicle trajectory (in the inertial frame), forces (in the body frame), kinematics (in the body frame), and the moments (in the body frame) [235]. For multi-copters such as VTOL, the forces include: aerodynamic drags caused by vehicle's motion and wind, motor-generated forces that are always in z-direction of the body frame, and the gravity force, which aligns with the z-direction of the inertial frame [235].

The energy consumption model contains two different methods, depending on the fidelity required: (1) simple, in which the power is the product of torque multiplied by the rotational speed, or (2) complex, in which DC motor power consumption is calculated as a product of the input voltage and current. For a given motor, both models show that the consumed energy is primarily a function of the motor's rotational speed, which is driven by the desired thrust and moments [235].

The collision avoidance model takes in the aircraft's position, speed, and intent, and intruder's position, speed, and intent, although some many require information from multiple intruders. FE<sup>3</sup> uses a trajectory projection based method, which is one that predicts the intruder's trajectory based on the intruder's current states (such as position and velocity) and ranks resolutions with predefined maneuver options and rules [235]. Models included are Detect and Avoid Alerting Logic for Unmanned Systems (DAIDALUS) and the Generic Resolution Advisor and Conflict Evaluator (GRACE) [235].

The communication and sensor models are restricted by the vehicle power and size. High transmission power increases the communication channel load and the signal interference, and essentially reduces the signal reception probability [235]. Since the high-density operation leaves a limited space for vehicles to avoid the potential conflicts, and the accuracy and update frequency of the intruder's states greatly affect the conflict avoidance algorithm's performance, the communication and detection capability plays a critical role in the analysis of the UTM-type traffic system [235].

The wind model used by FE<sup>3</sup> is The High Resolution Rapid Refresh (HRRR) system, developed by the National Oceanic and Atmospheric Administration (NOAA) [235]. This provides wind information at low-altitude airspace with great spatial and temporal resolutions. The HRRR wind data cover two altitudes at 10m and 80m [235]. The spatial resolution is 3 km by 3 km and its temporal resolution is 15 minutes [235]. When implemented in FE<sup>3</sup>, this model uses a spatially discretized database with turbulence intensity/uncertainty associated with every location [235].

## Calibration and Validation

Since FE<sup>3</sup>'s trajectory and energy consumption models are governed by first-principles physical equations, their results are well-understood and reasonable. In addition, FE<sup>3</sup>'s models for collision avoidance, communication, and wind, such as the DAIDALUS, GRACE, and HRRR, are well-known and fully calibrated and validated [235].

## 6.5 Simulation

The SimUAM toolchain is now applied to the San Francisco Bay Area as a case study. This toolchain combines network and origin-destination trip generation, ground traffic microsimulation, vertiport simulation, aerial traffic simulation, and UAM trip assignment and modeling. Once these steps are complete, the final step is to simulate both UAM and non-UAM trips[240, 235].

The first step is a standard MANTA simulation without UAM as a mobility option and with all trips as driving-only trips. These results are used for later comparison.

The next step involves running a MANTA simulation with UAM as a mobility option. In this scenario, there are several sequential steps. The first requires simulating in MANTA the access leg of each multi-modal UAM trip, with the rest of the trips in the demand assigned to driving.

After the toolchain runs, two experiments are presented, (1)  $S_1$ , Simulated road traffic in MANTA, and static vertiport transfer times and aerial flight times for the multi-modal UAM trips, and (2)  $D_1$ , Simulated road traffic in MANTA, simulated vertiport transfer times in VertiSim, and simulated aerial flight times in FE<sup>3</sup>.

In the static experiment,  $S_1$ , Once we have these initial times, we then simulate with static transfer and aerial times. A 2 minute transfer time is assumed at the origin vertiport. Then, each traveler is assumed to have access to an aircraft after 2 minutes at the origin vertiport, and they embark on their flight. This flight also has static characteristics, specifically 138 mph (120 kts) cruise speed, 2000 feet cruise altitude, ascent/descent angles of 10 degrees, and 100 fpm vertical climb. The flight path assumes a Haversine distance between the origin and destination vertiports. After the flight lands at the destination vertiport, a 2 minute transfer time at the destination vertiport is used.

In the dynamic experiment,  $D_1$ , instead of assuming static transfer and aerial times, travelers are simulated within the vertiport with VertiSim and are consolidated into aircraft. These aircraft's trajectories are then simulated in FE<sup>3</sup>. Once the flight lands at the destination vertiport, VertiSim, simulates the traveler's movement from the aircraft to their egress vehicle, as shown in Figure 6.1.

In both experiments, once the traveler reaches the egress vehicle, MANTA takes control and simulates the full demand with all access and egress trips for UAM as well as all the non-UAM driving trips.

## 6.6 Performance Metrics

SimUAM’s value is derived from its ability to produce a comprehensive set of performance metrics for all trips at both the microscopic person scale and the macroscopic regional scale. Once a network design is input into SimUAM, it is able to simulate all the demand in the region, assigning specific trips to multi-modal UAM and preserving the rest of the demand as driving trips.

In  $S_1$ , as mentioned earlier, the transfer times at the vertiports are statically assigned, and the aerial flights follow standard mission profile assumptions that depart on demand for each traveler. We evaluate the number of trips that benefit from taking multi-modal UAM, which is defined as a trip whose travel time when taking multi-modal UAM is shorter than the travel time of the same trip if it was driving only. Shown in Figure 6.8, we see that as the transfer times increase, the number of benefited trips decreases across all  $k$ . We can also see that the number of benefited trips increases as the number of vertiports increases, with the largest increase occurring between 10 and 30 vertiports, with more marginal benefit occurring as  $k$  increases further, ranging between 10K to 250K trips. This serves as a useful upper bound of the total addressable market of urban air mobility.




Figure 6.8: The number of benefited trips when the ground-air interface is not modeled

### Regional impact and addressable market

One of the major limitations of static assumptions is the inaccurate representation of aircraft being on-demand (i.e., being available to fly immediately when the traveler requires it). In experiment  $D_1$ , we now model the ground-air interface at the vertiport in which the wait times and consolidation times of passengers into aircraft, as well as the aircraft and vertiport capacities. Using the number of benefited from a transfer time of 2 minutes above, we now see after putting that demand through the full SimUAM toolchain that the number of benefited

trips at  $k = 10$ ,  $k = 30$ ,  $k = 50$ , and  $k = 70$  decrease significantly. In ??, we see that the 23K trips at  $k = 10$  to 5.3K trips, 125K trips at  $k = 30$  to 11K trips, 171K trips at  $k = 50$  to 18K trips, and 212K trips  $k = 70$  to 23K trips. This drop in the number of benefited immediately shows the impact of real world constraints that will exist for UAM at the vertiports themselves. Specifically, the number of takeoff and landing areas (TOLAs), number of parking spaces for the aircraft, and the number of passengers per aircraft will dictate the efficiency of multi-modal UAM as a beneficial transportation mode for travelers.  $k = 100$  is not calculated due to memory constraints for FE<sup>3</sup>.

In Figure 6.11, the number of benefited trips in the static  $S_1$  experiment shows a logarithmic growth over  $k$ , but once the ground-air interface is modeled in the loosely coupled SimUAM toolchain in  $D_1$ , we see that the number of benefited trips increases relatively linearly across  $k$ , showing that the waiting times and consolidation times at the vertiports significantly decreases the potential of multi-modal UAM for networks of 30 vertiports or fewer.




Figure 6.10: The number of benefited trips once the ground-air interface is simulated in SimUAM significantly decreases across all  $k$

The trips that benefit also change. Specifically, when considering the ground-air interface, the trips that benefit are those that have even higher travel times than when we did not consider the ground-air interface. For instance, at  $k = 10$ , the average driving time of trips that benefit when not considering the ground-air interface is 94 minutes, but when considering the ground-air interface, this increases to 101 minutes. The multi-modal UAM times for these same trips also increase when considering the ground-air interface, going from 64 minutes to 77 minutes. This is unsurprising given the increase in transfer times due to the ground-air interface bottleneck. This dynamic across  $k$  is shown in ??.




Figure 6.11: The average travel time of benefited trips once the ground-air interface is simulated in SimUAM significantly increases across all  $k$  (red vs. orange)

## Vertiports

Per the existing literature, we assume a throughput of 240 flights per hour, including both departures and arrivals [235]. At  $k = 10$ , the average transfer time once the ground-air interface is considered dramatically increases to 49 minutes, as shown in Figure 6.12. With a configuration of 4 TLOF pads, 16 parking spaces, and up to 4 seater aircraft, using the initial total demand of 23K at a transfer time of 2 minutes, we see that the real transfer times vary considerably across vertiports. The maximum average waiting time occurs at Vertiport 6, whose traffic of over 4K passengers forces the average transfer time to be over 2 hours. At Vertiport 4, since the demand is below 1K, the average transfer time is only 3.85 minutes.




Figure 6.12: The average transfer time at  $k = 10$  once the ground-air interface is considered

This shows that demand balancing may be important to decrease average waiting times

at each vertiport. However, balancing the demand in a network with few vertiports may also increase the access and egress driving times of these multi-modal UAM trips.

To investigate the effect of demand level on passenger transfer times, we have created 7 simulation scenarios with the vertiport configuration of 4 TLOF pads and 16 parking spaces per vertiport. 30 seconds of aircraft arrival per vertiport (arrival of 2 minutes per TLOF) is used for all cases. We use the same total demand, but now increase the number of vertiports in the UAM network. Thus, the demand per vertiport decreases. Figure 6.13 shows that as we decrease the demand, the average transfer time decreases down to an inflection point. The decrease in the average transfer time is due to the increased passenger inter-arrival time, which results in shorter queues for the flights. Decreasing the demand further actually results in an increase in the average transfer time because the accumulation time for the required number of passengers for a flight (4 passengers) increases. We also see that as we decrease the demand, the average aircraft occupancy decreases, due to the ratio of the passengers that wait more than  $T_{max}$  (10 minutes) increases and the flight generator is triggered for those passengers. Ultimately, waiting times decrease as  $k$  increases, but at 30 vertiports, the consolidation time increases as well, due to the demand being distributed across all vertiports rather than at certain vertiports.




Figure 6.13: The average transfer time and average aircraft occupancy with different demand levels

## Airspace

Once we simulate the aircraft in FE<sup>3</sup>, we see that there are now many aircraft congesting the airspace, with 7,352 flights for  $k = 30$  vertiports, shown in Figure 6.14. Note that these aircraft trajectories do not consider airspace restrictions, which will further affect the number of benefited trips if the airspace-restricted flight times are higher than the airspace-unrestricted flight times. FE<sup>3</sup> finds that with federated scheduling and no conflict resolution, the average delay is 21.5 seconds per flight.




Figure 6.14: The congested airspace of the San Francisco Bay Area at  $k = 30$  vertiports

## 6.7 Limitations and Future Research

The SimUAM toolchain has one notable limitation: the inability to do iterative convergence across a loosely coupled pipeline [YedavalliAssessing]. Due to the complex data input and output of iterative convergence, an ongoing solution is to cross-compile both GPU-based microsimulators, MANTA and FE<sup>3</sup>, and convert VertiSim into a GPU-based simulator. By synchronizing the time, the convergence issue will be eliminated.

Due to the long transfer times at the vertiports, future research should involve optimal network design to incorporate balanced vertiport networks. In addition, incorporating waiting time and consolidation time into the network design will be able to decrease the average transfer times across vertiports. Additional work involves refining the current components and adding new components. MANTA is currently being enhanced to carry out dynamic assignment, such that vehicles can adjust their routes based on the current congestion. FE<sup>3</sup> is currently being refactored to account for larger demand scales and vertiport networks, such as at  $k = 100$ , and also to consider local airspace restrictions. In addition, VertiSim is being restructured as a parallel programming architecture.

## 6.8 Conclusion

This paper proposes a high-level multi-component architecture to integrate and model the effects of introducing an emerging mobility technology, urban air mobility, into the transportation network. Scholarship hitherto has tackled different components of this pipeline



independently; however, in this research, we develop a full toolchain, SimUAM, to model the ground traffic, vertiport, and aerial traffic with granular, agent-based, high-fidelity simulation. Initial results show that once the ground-air interface is modeled, the market for UAM decreases across all network designs relative to models with static assumptions about transfer times. However, we also find that improvements can be made to balance the demand and optimize the networks for transfer time, likely increasing the number of benefited trips closer to the total addressable market once more. SimUAM is the first of its kind to incorporate fast, high-fidelity, regional-scale microsimulation on the ground and in the air with detailed ground-air interface modeling at the vertiport, and can be used as a powerful decision support system for UAM stakeholders.

# Chapter 7

## Evaluating the Risks and Readiness of U.S. Cities for UAM

### 7.1 Introduction

Advanced air mobility (AAM) is a broad concept focusing on emerging aviation markets and use cases for on-demand aviation in urban, suburban, and rural communities. AAM includes local use cases of about a 50-mile radius in rural or urban areas, and intraregional use cases of up to a few hundred miles that occur within or between urban and rural areas. Urban air mobility (UAM), a subset of AAM, envisions safe, sustainable, affordable, and accessible air transportation for passenger mobility, cargo delivery, and emergency management within or traversing a metropolitan area. In recent years, developments in electric and hydrogen propulsion; vertical and short takeoff and landing (VTOL/STOL) aircraft; and other technologies are creating renewed interest in on-demand urban aviation. A number of pre-COVID market studies estimate the market potential of UAM passenger mobility and goods delivery services to be between \$74 to \$641 billion US in 2035, with a wide variation due to study assumptions (e.g., geography, use cases, aircraft types, inclusion of military applications, etc.). Other studies also estimate a market potential for on-demand passenger services between \$2.8 to \$4 billion US in 2030 [176, 54, 9, 201].

However, introducing a new mobility mode into a regional transportation network does not occur in a black box. A new mode can create downstream effects - positive, negative, or both - as has been shown with the automobile over the twentieth century [199, 77, 22, 243]. The extent of these short-term and long-term effects depend on the geographic context, population investigated, type of deployment, and/or metric analyzed. Since transportation is a public good, planning the responsible integration of UAM, as with other modes, requires public stakeholder buy-in, both from agencies and their communities. With UAM, the public sector must interface with infrastructure developers, aircraft manufacturers, utilities, service providers, and communities to develop optimal policies and infrastructure for their respective constituents.

Given the vast number of planned deployments, federal preemption in many aspects of aviation regulation, and range of potential impacts UAM could have on communities, local and regional governments are faced with a situation where they need to evaluate the potential impacts of UAM while simultaneously planning for the potential growth of on-demand air mobility [220, 243, 216]. More research and planning are needed to help local and regional governments plan and prepare in order to mitigate the potential adverse impacts and maximize the potential benefits of UAM. While advanced air mobility may be enabled by the convergence of many factors, several challenges, including community acceptance, safety, equity, planning and implementation, airspace management, and operations, and others could create barriers to the growth and mainstreaming of UAM [54].

The objective of this paper is to survey local, regional, and state governments on their understanding of, planning capabilities, and institutional readiness for UAM. This paper is organized into four sections: First, the authors provide an overview of the methodological approach used to survey public agencies and an overview of the respondent demographics. In the next section, the authors evaluate the results of the survey, organizing the findings into three buckets: (1) broad UAM planning and concerns, with special attention given to safety, sustainability, equity, noise, and multi-modal integration, (2) stakeholder engagement, and (3) community engagement. In the final section, the authors conclude with limitations and recommendations for additional research.

## 7.2 Survey Methodology

This paper presents the results of a survey of public agencies across the U.S. to gauge familiarity, perceptions, and institutional readiness for UAM. The survey instrument was developed in partnership with the Community Air Mobility Initiative (CAMI), a non-profit established to educate and provide resources to the public and local and state decisionmakers. Study advisory members from CAMI, National Aeronautics and Space Administration (NASA), and faculty from UC Berkeley reviewed and commented on survey drafts, and completed a pretest of the survey.

The survey was organized into five key sections. The survey first asked respondents to answer questions about their agency's existing experience with aviation planning (if any), concerns with aviation impacts, and familiarity with UAM. Next, the survey asked questions about each agency's planning process and if they currently or are considering planning for UAM. The survey also asked questions about the potential UAM use cases that may be under consideration by public agencies. Additionally, the survey asked questions about existing and planned stakeholder and community engagement activities with respect to UAM. The final section of the survey explores specific UAM planning concerns, focusing on safety, sustainability, equity, noise, and multimodal integration.

The survey was conducted online using the Qualtrics platform from February to June 2021. Using publicly available databases, the authors assembled contact information for key personnel at cities, metropolitan planning organizations (MPOs), local emergency man-

agement agencies, state departments of transportation (DOTs), and public transit agencies from the largest 100 metropolitan statistical areas (MSAs) in the U.S., and the three largest MSAs in each state. The authors emailed a short description and a link to take the survey. C-level officials, division or department heads, practitioners and planners, and elected officials responded to the survey. Approximately 500 people were contacted, yielding 105 responses, or a 21% response rate. However, respondents were not required to answer each question and the response rate for some questions may be lower.

A summary of the institutional respondents and their role within their organization is shown in Figure 7.1. Out of the 105 responses obtained, 47% of the sample represent MPOs, 24% representing cities, 16% represent state DOTs, 11% represent public transit agencies, and the remainder represent other entities, such as emergency management agencies. Additionally, the sample includes a sizable percentage of C-suite respondents (17%), division directors (41%), and practitioners (32%). A map of respondents by location and jurisdiction type is shown in Figure 7.2.




Figure 7.1: Respondents by jurisdiction type and by role within their organization

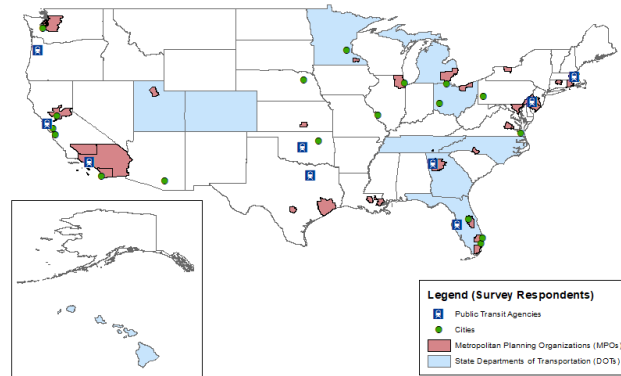


Figure 7.2: Respondents by jurisdiction type and location across the United States

## 7.3 Survey Findings

The survey explores a variety of issues relevant to government agencies when planning for UAM. This section is organized into two sections: (1) planning for UAM, with particular detail given to the issues of safety, sustainability, equity, and noise, and (2) UAM integration in the transportation network.

### Planning for UAM

The planning process allows planners and policymakers to document the state of transportation networks (e.g., congestion and performance), and establish goals and policies to guide future growth and infrastructure development (Cohen and Shaheen 2016). Integrating UAM into the planning process can serve a few core functions. First, the potential impacts of UAM can be modeled and forecasted as part of long-range planning efforts. Second, planners can integrate UAM as part of infrastructure capital and planning decisions, such as where vertiports get placed and how they are integrated with surface transportation. Finally, integrating UAM into the planning process can allow public agencies to study the impacts of UAM airspace corridors and vertiport placement on land use, circulation, noise, safety, and other common elements of comprehensive plans.

In the U.S., federal preemption in aviation gives the Federal Aviation Administration (FAA) exclusive authority to regulate aviation safety, the efficiency of the navigable airspace and air traffic control, among other things [170]. For this reason, local, regional, and state governments are not permitted to regulate aircraft operations, such as flight paths, altitudes, or navigable airspace. However, local, regional, and state governments retain wide latitude over the regulation of take-off and landing sites through land use and zoning. Typically, laws related to local policy, such as land use, zoning, privacy, and law enforcement operations, are generally not subject to federal regulation [170]. For these reasons, aviation stakehold-

ers (e.g., air carriers, manufacturers, etc.) have generally had limited interaction with local stakeholders, such as elected officials, planning departments, public works, and underserved communities. These relationships are depicted in Figure 7.3. However, because UAM decentralizes the impacts of aviation beyond airports, understanding the potential impacts of on-demand air mobility and the role of local and regional governments is important.

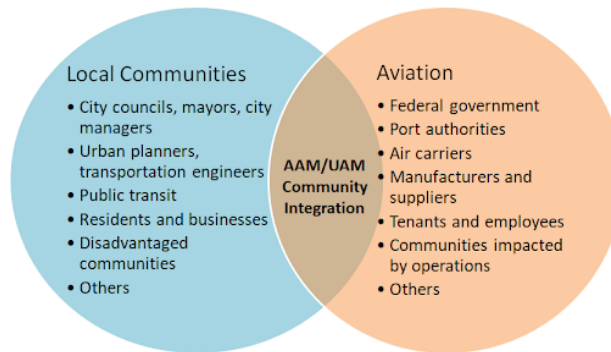


Figure 7.3: The Relationship between local communities, aviation stakeholders, and UAM community integration

Recognizing this historical context, the survey first collected baseline data on the existing aviation activities (if any) and familiarity with UAM. The survey asked three key questions. First, public agencies were asked if they currently operate any aviation facilities, such as airports or heliports. Second, public agencies were asked if they currently incorporate UAM as part of existing planning documents. Third, public agencies were asked if they were considering adding UAM into future planning documents.

### Experience with Aviation and Incorporating UAM into Local, Regional, and State Planning Processes

A number of local, regional, and state governments are engaging in a variety of planning activities to prepare for UAM. In May 2021, the Texas senate passed Senate Bill 763 which creates an Urban Air Mobility Advisory Committee at the Texas Department of Transportation to assess current and potential state laws and regulations needed to facilitate the development of UAM operations and infrastructure [205]. In May 2021, the Massachusetts, Minnesota, and Ohio Departments of Transportation, as well as the North Central Texas Council of Governments and the City of Orlando have signed agreements with NASA to participate in a workshop series to identify practices that states and municipalities can use to integrate UAM into local and regional transportation plans [95]. The City of Orlando has developed a White Paper establishing a preliminary framework to explore how UAM

could be integrated into the city's transportation network [161]. In 2020, Los Angeles and the World Economic Forum published Principles of the Urban Sky, which outlines seven planning tenets necessary to ensure the sustainable and equitable integration of UAM into cities [171]. These principles include planning and implementation considerations, such as safety, sustainability, equity, and multimodal integration, among others.

The survey asked respondents about their existing experience with aviation, and existing and planned efforts to incorporate UAM into local, regional, and state planning processes. Among the public agencies responding to the survey, only a small percentage (9%) of the sample had any direct experience operating aviation facilities, such as airports and heliports. Seventy-eight percent of the sample that operated aviation facilities were municipal governments. Forty-seven percent ( $n=49/105$ ) of the public agencies responding to the survey indicated that they are currently incorporating UAM into existing plans. The public agencies that most frequently reported incorporating UAM into existing plans were MPOs (53% or  $n=26/49$ ) and cities (44% or  $n=11/25$ ). Comprehensive plans—also known as general or master plans—are plans with a set of long-term goals and policies that communities use to guide development decisions (Cohen and Shaheen 2016). Twelve percent of municipal governments responding to the survey are incorporating UAM into comprehensive plans. Additionally, 32% ( $n=8/25$ ) are considering incorporating UAM into comprehensive plans and 20% ( $n=5/25$ ) into subarea plans. Similar to comprehensive plans, regional transportation plans (RTPs) are long-range plans that identify and analyze the transportation needs of a metropolitan region. Twelve percent of MPOs responding to the survey are already incorporating UAM into these RTPs ( $n=6/49$ ). An additional 25% of MPOs ( $n=12/49$ ) are considering incorporating UAM into planning documents, such as regional transportation plans (20%) and comprehensive plans (14%).

Eighty-two percent of cities and 64% of MPOs indicated that they are not currently or considering incorporating UAM into their plans. Lack of familiarity with UAM or with aviation planning were the two most important reasons that cities and MPOs are not currently planning for UAM. Sixteen percent of cities and 27% of MPOs said they lacked familiarity with UAM. Sixteen percent of cities and 22% of MPOs said they lacked familiarity with aviation planning, which makes it difficult to incorporate UAM into their planning processes. These findings suggest that more education and outreach is needed to help cities and MPOs understand what UAM is, and how to plan for it, particularly given the lack of familiarity with aviation among local and regional governments. Section 7.3 reflects these findings across the different jurisdiction types.

As an emerging concept, UAM will face many barriers, particularly at the local and regional levels of government, such as community concerns about noise, equity, multimodal integration, safety, and sustainability. The survey asked respondents about their agency's perceptions toward UAM. Respondents were asked to rate their agency's concerns about UAM using a 10-point Likert scale ranging from 1 (not concerned at all) to 10 (extremely concerned). The question consisted of 15 types of concerns commonly associated with UAM, including those related to safety, sustainability, privacy, economic productivity, affordability, neighborhood impacts, and visual pollution, among others. In doing so, the survey provides

Jurisdiction Type	% Operating Existing Aviation Facilities	% Incorporating UAM into Existing Planning Documents	% Considering Incorporating UAM into Future Planning Documents
City (n=25)	7% (n=7)	10% (n=11)	8% (n=8)
MPO (n=49)	0% (n=0)	25% (n=26)	11% (n=12)
State (n=17)	1% (n=1)	7% (n=7)	5% (n=5)
Public transit agency (n=11)	0% (n=0)	3% (n=3)	3% (n=3)
Other (n=3)	1% (n=1)	2% (n=2)	2% (n=2)

Table 7.1: UAM Planning by jurisdiction type (Notes: (1) Percentages per column will not add up to 100% due to how the data is presented. Instead, the sum of percentages in each column will signify the total percentage of respondents who answered Yes to the question, with the remaining percentage as No or Unsure), (2) Each percentage is rounded to the nearest whole number.

early insights about different types of concerns among different types of public agencies (e.g., city, MPO, and state), and comparative levels of concern about various impacts among a cohort of public agencies (e.g., municipal concerns about safety compared to other types of UAM impacts). The survey found that noise was the most important concern among city governments, followed by equity. In contrast, equity (followed by safety) were the most important concerns of MPOs and state DOTs.

**Noise** Noise from surface transportation, aviation, and emergency services are frequently cited as a common community concern in urban environments. In some cases, communities have responded to concerns about aviation noise by closing airports and heliports. For example, the City of Santa Monica is scheduled to close its airport in 2028. A number of aspects of UAM could create noise nuisances for communities. Potential nuisances include how loud an aircraft’s noise is, how an aircraft’s noise sounds, how long the noise persists, the types of land uses where the noise is heard, the time of day the noise occurs, and differences in noise characteristics associated with scaled operations (both number of aircraft and frequency of flights). Noise was the highest ranked concern among municipal governments. Twenty-eight percent of cities (n=7/25) reported having an ordinance restricting vehicular or truck noise. Six percent of MPOs (n=3/49) also reported having similar restrictions on vehicle noise. Generally, cities and MPOs had similar views on noise. Both were most concerned



by the volume of the noise, followed by the time of day the noise occurs. In contrast, state DOTs were more concerned with how the noise sounds (i.e., frequency/type of noise) and secondarily the volume of UAM noise. Due to concerns about aircraft noise, UAM aircraft may need to meet stricter noise standards due to anticipated low-level flight paths over densely populated areas.

**Equity** Social equity may be one of the largest barriers to the growth of UAM. Common equity concerns associated with UAM include the high-cost of flights that may not be affordable to a mass market; environmental impacts of flight paths and vertiports, such as noise, traffic (air and surface), and visual pollution; and economic impacts, such as gentrification and displacement, in neighborhoods around vertiports. Additionally, limited public funds and expenditures on UAM could result in other foregone public investments that could have a broader societal impact, such as investments in public transportation. Current on-demand aviation services average approximately 149to300 US per seat (although some are considerably more expensive). Today, these price points serve higher-income and business travelers, and are inaccessible for the mass market. A study by Porsche Consulting forecast that on-demand air taxis using eVTOL will cost \$8 to 18 US per minute [201]. Another study by McKinsey and Company estimates that UAM air taxi service will range between \$131 to \$1,912 US per trip in 2030 (depending on vertiport density) [142].

MPO and state DOT respondents rated equity as their agency's most important concern associated with UAM. Equity was also ranked as the second most important concern among municipal governments. For local governments and MPOs, the affordability of UAM and the environmental impacts of both flight paths and vertiports were identified as notable concerns. For state DOTs, the affordability of UAM and the economic impacts of vertiports on neighborhoods, such as gentrification and displacement, were key concerns. Vertiport and aircraft access for people with disabilities was also raised as a somewhat important equity concern among all three agency types. The findings indicate that policies are needed to expand affordability of UAM, ensure access for people with disabilities, and prevent or mitigate the environmental and economic impacts of vertiports and UAM operations. Across all jurisdiction types, affordability is the most important priority for agencies across all jurisdiction types, with an average priority level of 5.0 out of 10. Environmental justice and ADA access are next, with average ranks of 4.38 and 3.65, respectively. Economic impacts on neighborhoods and social justice are prioritized in the middle, and those with less access to services and tools such as banks and phones are prioritized among the lowest. Job creation is the least concerning issue when evaluating the equity of emerging transportation technologies, as shown in Figure 7.4.




Figure 7.4: Average priority level of equity concerns across all jurisdiction types. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned.

**Multimodal Integration** The survey also asked if public agencies would be interested in integrating UAM with existing surface transportation options. Both cities and MPOs indicated that public transit and shared mobility were the most important modes of transportation for UAM to connect with. State DOTs, however, indicated that it was most important to connect UAM to shared mobility and private vehicles. This variation likely reflects differences in roles and responsibilities, as many state DOTs are generally more focused on roads and highways. Twenty four percent of state DOTs that responded to the survey ( $n=4/17$ ) said they were considering integrating UAM with public transit. Additionally, 20% of both cities and MPOs ( $n=5/25$  and  $n=10/49$ , respectively) also said their agency was considering integrating UAM with public transportation. One MPO and one state DOT (2% and 6%, respectively) also said their agencies were considering integrating UAM with shared mobility. Another two MPOs and one state DOT (4% and 6%, respectively) said they were planning on integrating UAM with a mobility-as-a-service platform where travelers can access multiple transportation modes over a single digital interface. Notably, across all jurisdiction types, the least important mode to integrate with was any version of private vehicle, including privately-owned vehicles and privately-owned automated vehicles, shown in Figure 7.7.




Figure 7.5: Average priority level of integrating with existing modes across all jurisdiction types. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned.

**Safety** Supplemental expert interviews on the topic revealed that public agencies were generally less concerned about safety due to the aviation sector’s safety record. However, the interviewees were also unfamiliar with the different safety records of Part 91, Part 135 (scheduled), Part 135 (non-scheduled), and Part 121 services. These parts refer to the U.S. Federal Aviation Regulations (FARs) that govern all aviation activities in the U.S. Part 91 of the FARs provides general operating and flight rules for small non-commercial aircraft within the U.S. Part 135 primarily oversees commuter and on-demand operations. Part 121 regulates commercial airlines. While each of these service models have very different safety records, the public sector generally perceived most aviation activities generally safe due to the safety record of Part 121 carriers (i.e., airlines). This suggests that public agencies have a high-level of trust in the aviation sector and believe that UAM services will have an equal or better safety record than commercial airlines. More education and outreach may also be needed to educate local and regional governments on aviation safety, and the factors that contribute to variations among different types of FAA-regulated activities. The findings confirm this, as cities and MPOs indicate that the number of fatalities is the highest priority concern for UAM, but with an average priority level of 7.24 out of 10.0. However, states indicate that the number of injuries is their highest priority concern, with an average priority level of only 7.20 out of 10.0. Property damage is indicated as the least concerning safety issue across all jurisdictions. Figure 2 gives further detail.




Figure 7.6: Average priority level of safety concerns across cities, states, and MPOs. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned.

**Sustainability** Investments in UAM are closely linked to recent technological developments in electric, hydrogen, and other low and zero emission aircraft. Proponents of UAM suggest a variety of potential environmental benefits, such as shifting single-occupant vehicle drivers into shared electric aircraft with multiple passengers. However, UAM also has the potential to induce demand due to factors such as reduced travel times. While more research is needed to study the potential environmental impacts of UAM on communities, this survey asked respondents to rate their environmental priorities with respect to the introduction of UAM as a new mode of transportation. The survey found that reducing greenhouse gas emissions were the most important UAM environmental priorities for municipalities. However, improving air quality was most important to MPOs and state departments. Despite the variation in the top priority, reducing vehicle miles traveled (VMT) was ranked as the second-highest priority across all three jurisdiction types.




Figure 7.7: Average priority level of sustainability concerns across cities, states, and MPOs. The scale is out of 10.0, with 10.0 being most concerned and 1.0 being least concerned.

### The Role of Stakeholder Engagement

As the public sector prepares for early UAM deployments, both collaboration with UAM stakeholders and community engagement will be important parts of the planning process to help communities understand potential concerns, such as noise, aesthetics, traffic congestion, equity (i.e., displacement and gentrification), safety, and others. Respondents were asked to share information about their collaboration with UAM stakeholders. Twenty-eight percent of cities (n=7), 27% of MPOs (n=13), and 29% of state respondents (n=5) reported currently collaborating with key UAM stakeholders, such as organizations involved in education and outreach; service providers; aircraft manufacturers; and vertiport and real estate developers. Municipal governments most frequently reported collaborating with organizations involved in education and outreach (24%, n=6/25), service providers planning on using novel aircraft designs such as eVTOL (24%, n=6/25), and aircraft manufacturers (16%, n=4/25). A few municipal respondents also reported collaborating with utility companies (12%, n=3/25), vertiport developers (4%, n=1/25), real estate developers (4%, n=1/25), and UAM service providers currently offering on-demand flights using helicopters (4%, n=1/25). Similarly, MPOs also most frequently reported collaborating with organizations involved in education and outreach (18%, n=9/49); service providers planning on using novel aircraft designs (12%, n=6/49); and aircraft manufacturers (12%, n=6/49). One MPO (2%) also reported working with a utility company on electric infrastructure for UAM.

Of the local and regional governments collaborating with UAM stakeholders, the majority reported discussions related to planned passenger services using on-demand air taxis. While local and regional governments more frequently reported ongoing collaboration with organizations involved in education and outreach, aircraft manufacturers, and service providers, state respondents reported more varied collaboration activities. Three states (18%) reported collaborating with organizations involved in education and outreach. Two states (12%) said they were collaborating with UAM service providers using novel aircraft designs, utility com-

panies, and UAM service providers using helicopters. One state (6%) reported collaborating with aircraft manufacturers, vertiport developers, and other real estate developers. The results indicate a variety of emerging planning activities related to UAM, even though UAM has not been widely integrated into local, regional, and state planning documents.

### The Role of Community Engagement

Community engagement can play a crucial role in helping prepare the public for early deployments of UAM and guide sustainable and equitable outcomes. In spite of relatively robust stakeholder engagement, survey respondents reported much more limited community engagement efforts. Only two cities (8%) and two MPOs (8%) reported currently working with community-based organizations as part of UAM outreach. One city (4%) and three MPOs (6%) said they were engaging local business leaders. Similarly, only two cities (8%) indicated they were working with neighborhood leaders and associations on UAM-related outreach activities. The lack of early community engagement with community stakeholders could be due to the state of the industry of UAM at the time of the survey, such as limited program deployments, lack of certified eVTOL aircraft, and the global pandemic, which may have limited the ability for public agencies to engage their communities through in-person methods. However, the lack of community engagement could also present a number of risks for the public and private sectors if communities impacted by UAM are not provided opportunities for meaningful engagement in the transportation decision-making process. These results suggest that the public and private sectors should enhance community engagement activities around UAM planning and implementation with key constituencies, such as business leaders, neighborhood associations, community-based organizations, residents, and others.

A variety of methods can be used to solicit public feedback as part of the UAM planning process, such as in-person and virtual public meetings, workshops, focus groups, listening sessions, social media, and websites. The survey also asked respondents to identify the community engagement strategies their agency thought were the most suitable for UAM planning and community engagement. Municipal governments indicated that the most effective community engagement strategies for UAM are focus groups (40%,  $n=10/25$ ), social media (36%,  $n=9/25$ ), and listening sessions with the public (32%,  $n=8/25$ ). MPOs reported that the most effective UAM community engagement strategies are public meetings (35%,  $n=17/49$ ) and surveys (35%,  $n=17/49$ ).

## 7.4 Limitations and Additional Research

There are three major limitations associated with this study. The first is sample size and uniformity across the different types of jurisdictions, regions, and MSA populations. Our sample is heavily skewed toward large metropolitan areas, primarily on the Pacific Coast, Midwest, and Mid-Atlantic regions. A larger sample of responses across geography and jurisdiction type could enable a more granular analysis.

Another limitation is the survey's focus on only noise as part of community acceptance. While results show that many agencies have not yet formally engaged with their communities, it is likely that issues outside of noise will also be important to them, such as affordability and aesthetics. Future research should consider cross-referencing community concerns with those of the agencies.

The final limitation is associated with the survey design. By not forcing responses to particular questions, many agencies did not answer the entire survey, making cross-tabulation for particular parameters challenging. An enhancement for future survey instruments would be to force responses for all questions, acknowledging that the response rate may decrease as a result.

Additional research includes using these data to understand next steps for targeted policymaking for the large MSAs, which have already shown greater preparedness and engagement in UAM. A monolithic approach to regulation will not suffice, as regional variation and regulatory differences for each jurisdiction type requires bespoke policy.

## 7.5 Conclusion

Over 100 public agencies provided detail about their engagement with UAM. As the first comprehensive analysis of U.S. public agencies' preparedness for UAM, this study revealed many new findings. Planning priorities, the level of stakeholder engagement, and the type of community engagement vary across jurisdiction type. All jurisdiction types trust that safety will follow the rigorous standards set forth by the FAA for commercial aviation. All jurisdiction types also find that integrating UAM with public transit, shared modes, and active modes is important, while integrating with private vehicles is a much lower priority. Variation across jurisdiction type exists in sustainability, noise, and equity. Stakeholder engagement has been relatively high across all jurisdiction types, with variation across jurisdiction type among which stakeholders have been engaged. However, community engagement remains low across the board and must be highlighted as a top priority for public agencies in the near future.

The Pacific Coast is especially advanced in their current engagement with aviation as well as their engagement with UAM in particular, with the Midwest and Southwest not far behind. Large MSAs, particularly in the Pacific Coast and the southern United States, have been preparing for UAM. Medium- and small-sized MSAs, on the other hand, are not only not thinking about UAM, but are often not aware of or familiar with it. Accordingly, UAM deployments are expected to be in the larger MSAs initially, likely beginning with goods transport and slowly moving toward passenger mobility. These findings suggest that targeted regional policymaking and infrastructure efforts for UAM is necessary to move UAM forward.

# Chapter 8

## Conclusion

### 8.1 Summary of Contributions

The twentieth century saw the rise of the automobile, and with it came dramatic shifts in travel behavior and infrastructure provision. Over time, the increased reliance on the automobile has caused devastating gridlock, higher pollution, and unsustainable land use and travel behavior in metropolitan areas across the world. Considering the significant implications to sustainability, equity, and productivity, transportation planning has become more important than ever. In order to avoid the downstream effects of new mobility technologies, greater care must be taken with their deployments. Now, planners, engineers, and policy-makers are looking for alternatives to battle the growing transportation problem. Broad types of initiatives exist, such as providing incentives for behavioral changes in mobility (e.g. congestion pricing), replacing trips (e.g. telecommuting), and enhancing modal substitution by offering a wide variety of accessible transportation alternatives. Urban air mobility (UAM) is emerging as one of these alternatives.

Introducing a new mode must go through a comprehensive transportation planning process to understand its short- and long-term impacts. As it stands, the impacts of UAM, both positive and negative, remain unknown. This dissertation broadly explores urban air mobility by developing a software framework to understand its potential impacts and then investigating government readiness to deploy UAM in their areas.

A novel traffic microsimulator, MANTA, which addresses the challenges of fast, accurate traffic microsimulation at the metropolitan-scale, is first presented. MANTA is highly efficient and is capable of simulating real-world traffic demand with a fine level of granularity on very large-scale networks, offering significant potential for fast scenario planning in both short- and long-term applications in metropolitan-scale analysis. Independent functionality to infer traffic control using deep learning is then discussed, with future plans to incorporate it into MANTA. The multi-layer convolutional neural network produces the accuracies of approximately 72%, classifying the correct traffic control in 3 out of 4 intersections, significantly improving the state-of-the-art.



UAM as a transportation mode is then integrated into the traffic microsimulation architecture. UAM's benefit in a metropolitan area is first understood under specific land use constraints for vertiports. Its value is then analyzed at a more granular level using an iterative, data-driven approach to determine the total addressable market for UAM. These demands and networks are simulated in MANTA as multi-modal UAM trips along with driving trips across the entire metropolitan area. Results show that short distance trips are gained by expanding the network to a higher  $k$ , and the long distance trips benefit across all network designs. A notable increase in market potential exists between  $k = 10$  and  $k = 30$  vertiports. Additionally, at all numbers of vertiports, the total regional congestion with UAM as a mode is lower than if there was no UAM and all trips were driving trips. Finally, non-UAM users are shown to benefit very marginally across all  $k$ , with up to 1 minute in benefit on average at  $k \geq 30$ .

These initial analyses assume static transfer times at the vertiports and basic aerial flight time calculations, but better real-world modeling is necessary to understand the ground-air interface at the vertiport. As a result, a holistic software toolchain, SimUAM, which integrates high-fidelity ground traffic simulation, vertiport simulation, and aerial simulation, is developed. Rather than assuming the ground-air interface as static, this analysis reveals the real-world bottlenecks in UAM network optimization: vertiport capacities and aircraft capacities, reflected in large waiting and consolidation times. Considering these constraints, average transfer times fluctuate wildly across vertiports in different network designs, making UAM load balancing an important future research topic and significant in dictating the viability of UAM at scale.

The dissertation concludes with an evaluation of the preparedness of government agencies for UAM deployment. Results show that large metropolitan statistical areas (MSAs), particularly in the Pacific Coast, Midwest, and Southeast, have been actively engaging with stakeholders in UAM and have already incorporated or are planning to incorporate UAM into their plans. In addition, the majority of metropolitan planning organizations, states, and cities are concerned about UAM's equity implications. While many agencies have broad plans for UAM, they have engaged minimally with their respective communities, suggesting that public acceptance may become a bottleneck to UAM deployment. Hurdles for public agencies exist, but these results indicate that UAM is no longer a futuristic mobility concept, but instead a real transportation alternative that government agencies of many large MSAs are already planning for. As a result, this interest further underscores the need for transportation modeling tools such as SimUAM as decision support systems.

## 8.2 The (Sky) Ahead...

The simulation toolchain developed in this dissertation is necessary, but not sufficient. Even deeper analysis and scenario planning must be conducted to understand UAM's broader impacts.

One of the main impacts that must be analyzed is the interaction between travel behavior

and long-term land use. Most existing UAM studies assume that demand is exogenous to network design. However, in reality, potential changes in travel time may lead to changes in individual behavior altogether, which then has a significant impact on residential and workplace location choices, and ultimately land use, in the long run. This impact is contingent on the placement of vertiports to not only maximize the number of travelers, but also to decrease the access and egress times to the vertiports themselves. Urban economic theory suggests that both centrality and dispersion may result from a new mode like UAM, and thus planning UAM with particular objectives of sustainability and equity in mind is imperative [6, 216, 36]. The SimUAM toolchain has been made modular and extensible to enable this functionality, as shown in Figure 8.1.

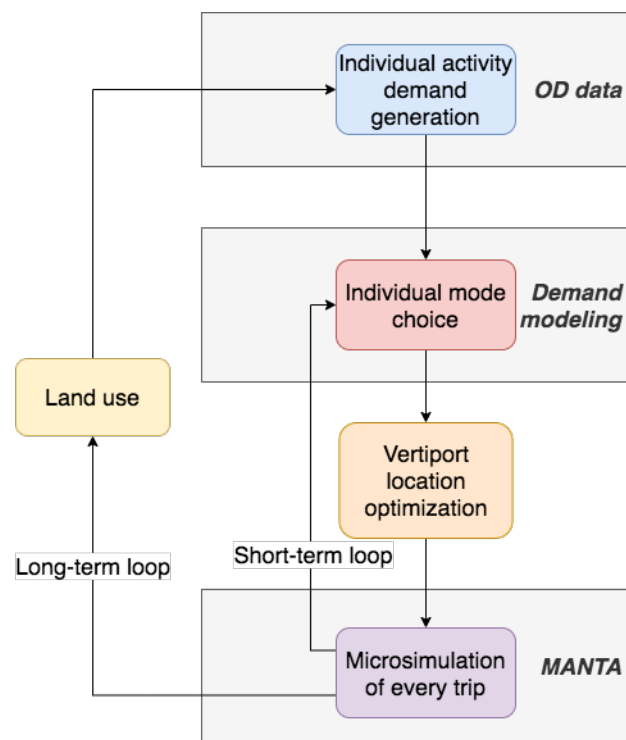


Figure 8.1: Incorporating the long-term land use loop

Another challenge is developing beneficial UAM network designs that optimize for safety. Infrastructure for vertiports, vehicle charging, and airspace traffic management must mature to maximize safety [77]. However, a feedback loop exists - as more users take UAM for daily trips, congestion will begin to appear in the sky, which will not only cause concern for the riders but also increase safety risks for both riders and the overflow [58, 155].

A related obstacle is the rate of adoption. From the rider perspective, Garrow et al. show that stated ridership is much greater within younger, more affluent communities, and anticipate that the first application will be time-sensitive business trips [90]. In addition,

despite high capital costs associated with private ownership, studies show that Americans prefer the personal vehicle model of flying cars over ridesharing options, which may limit the scale of adoption, and accordingly may decrease any positive impacts in sustainability and efficiency [155].

Economic productivity is yet another unknown impact of UAM. If travel time decreases altogether, more time can be spent either at work or in leisure activities that may spur the local economy [23]. However, in this scenario, induced demand may also create greater congestion on the road and even in the skies. On the other hand, scenarios may exist where UAM increases total travel times, rendering it an inadequate transportation alternative. In preliminary research with AV, there are indeed economic productivity gains from decreased travel times and a lower number of accidents, but ripple effects occur in which there is lower demand for auto repair, traffic services, medical, insurance, and legal, leading to \$1.2 trillion in economic impact [51]. Similar effects could occur with UAM, but these analyses must be adapted to every metropolitan or regional economy.

Finally, and perhaps most importantly, based on public agencies' responses, understanding the equity implications of UAM is a significant area of research. UAM's effect may be similar to the current impacts of ridesharing, which has often become a substitute for transit rather than single-person automobile, and thus, careful planning must be done to ensure more equitable outcomes [52].

### 8.3 Final Thoughts

Urban air mobility has real opportunity to provide benefits to travelers in metropolitan areas. However, a fine line exists between benefit and harm. It is imperative that we understand how to responsibly deploy this new mode to optimize for sustainability, equity, and efficiency, among many other transportation planning objectives. As this dissertation has detailed, developing the proper decision support systems, which include accurate UAM models and extensible simulation architectures, is necessary before committing valuable resources and time to deployment.

UAM enthusiasts and policymakers alike are now deeply engaged, and experts in various disciplines from planning to transportation to computer science must be both visionary and responsible at the same time. Integrating technology into society has never been at a more significant inflection point, and understanding how to use technology to do service to society will be the key to future success.

# Bibliography

- [1] AC Transit, Arup, and Cambridge Systematics. *Bay Bridge Corridor Congestion Study*. TJPA, 2010.
- [2] Justin Guan Adam Cohen. “Reimagining the Future of Transportation with Personal Flight”. In: (2020). DOI: 10.7922/G2TT4P6H. URL: <https://escholarship.org/uc/item/9hs209r2> (visited on 05/13/2020).
- [3] Eric Adams. *Covid-19 Casts a Dark Cloud Over the Flying Car Future — WIRED*. WIRED. URL: <https://www.wired.com/story/coronavirus-covid-air-taxi-flying-car-impact/> (visited on 05/12/2020).
- [4] Adrian Albert, Jasleen Kaur, and Marta C. Gonzalez. “Using Convolutional Networks and Satellite Imagery to Identify Patterns in Urban Environments at a Large Scale”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Halifax, NS, Canada). KDD '17. New York, NY, USA: ACM, 2017, pp. 1357–1366. ISBN: 978-1-4503-4887-4. DOI: 10.1145/3097983.3098070. URL: <http://doi.acm.org/10.1145/3097983.3098070> (visited on 04/21/2019).
- [5] William Alonso. *Location and Land Use*. Feb. 5, 1964. URL: <https://www.degruyter.com/view/product/249205> (visited on 02/14/2019).
- [6] Alex Anas, Richard Arnott, and Kenneth A. Small. “Urban Spatial Structure”. In: *Journal of Economic Literature* 36.3 (1998), pp. 1426–1464. ISSN: 0022-0515. JSTOR: 2564805.
- [7] James Anderson et al. *Autonomous Vehicle Technology: A Guide for Policymakers*. RAND Corporation, 2016. ISBN: 978-0-8330-8398-2 978-0-8330-8437-8 978-0-8330-8438-5 978-0-8330-8439-2. DOI: 10.7249/RR443-2. URL: [http://www.rand.org/pubs/research\\_reports/RR443-2.html](http://www.rand.org/pubs/research_reports/RR443-2.html) (visited on 02/14/2019).
- [8] Kevin R. Antcliff, Mark D. Moore, and Kenneth H. Goodrich. “Silicon Valley as an Early Adopter for On-Demand Civil VTOL Operations”. In: *16th AIAA Aviation Technology, Integration, and Operations Conference*. 0 vols. AIAA AVIATION Forum. American Institute of Aeronautics and Astronautics, June 10, 2016. DOI: 10.2514/6.2016-3466. URL: <https://arc.aiaa.org/doi/10.2514/6.2016-3466> (visited on 02/22/2019).

- [9] *Are Flying Cars Preparing for Takeoff?* Morgan Stanley. URL: <https://www.morganstanley.com/ideas/autonomous-aircraft> (visited on 07/02/2021).
- [10] Stephen Arhin. *Predicting Acceptable Wait Times for Patrons at Transit Bus Stops by Time of Day*. July 2018. URL: <https://transweb.sjsu.edu/mctm/research/utc/Predicting-Acceptable-Wait-Times-Patrons-Transit-Bus-Stops-Time-Day> (visited on 10/30/2020).
- [11] Richard Arnott and Kenneth Small. “The Economics of Traffic Congestion”. In: *American Scientist* 82.5 (1994), pp. 446–455. ISSN: 00030996. JSTOR: 29775281.
- [12] Richard Arnott and Kenneth Small. “The Economics of Traffic Congestion Rush-Hour Driving Strategies That Maximize an Individual Driver’s Convenience May Contribute to Overall Congestion”. In: (). JSTOR: [pdf/29775281.pdf](https://www.jstor.org/stable/29775281).
- [13] Kai Arulkumaran et al. “A Brief Survey of Deep Reinforcement Learning”. In: *IEEE Signal Processing Magazine* 34.6 (Nov. 2017), pp. 26–38. ISSN: 1053-5888. DOI: 10.1109/MSP.2017.2743240. arXiv: 1708.05866. URL: <http://arxiv.org/abs/1708.05866> (visited on 08/15/2018).
- [14] Joshua Auld et al. “POLARIS: Agent-Based Modeling Framework Development and Implementation for Integrated Travel Demand and Network and Operations Simulations”. In: *Transportation Research Part C: Emerging Technologies* 64 (Mar. 2016), pp. 101–116. ISSN: 0968090X. DOI: 10.1016/j.trc.2015.07.017. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0968090X15002703> (visited on 05/16/2020).
- [15] Kay W. Axhausen and ETH Zürich. *The Multi-Agent Transport Simulation MATSim*. Ed. by ETH Zürich et al. Ubiquity Press, Aug. 10, 2016. ISBN: 978-1-909188-75-4. DOI: 10.5334/baw. URL: <http://www.ubiquitypress.com/site/books/10.5334/baw/> (visited on 05/01/2020).
- [16] Kay W. Axhausen and Tommy Gärling. “Activity-based Approaches to Travel Analysis: Conceptual Frameworks, Models, and Research Problems”. In: *Transport Reviews* 12.4 (Oct. 1, 1992), pp. 323–341. ISSN: 0144-1647. DOI: 10.1080/01441649208716826. URL: <https://doi.org/10.1080/01441649208716826> (visited on 04/19/2019).
- [17] Saeed Asadi Bagloee et al. “Autonomous Vehicles: Challenges, Opportunities, and Future Implications for Transportation Policies”. In: *Journal of Modern Transportation* 24.4 (Dec. 2016), pp. 284–303. DOI: 10.1007/s40534-016-0117-3. URL: <http://link.springer.com/10.1007/s40534-016-0117-3>.
- [18] Scott Baker et al. *How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic*. w26949. Cambridge, MA: National Bureau of Economic Research, Apr. 2020, w26949. DOI: 10.3386/w26949. URL: <http://www.nber.org/papers/w26949.pdf> (visited on 07/02/2021).

- [19] Milos Balac, Raoul L. Rothfeld, and Sebastian Hörl. “The Prospects of On-Demand Urban Air Mobility in Zurich, Switzerland”. In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. 2019 IEEE Intelligent Transportation Systems Conference (ITSC). Oct. 2019, pp. 906–913. DOI: 10.1109/ITSC.2019.8916972.
- [20] Milos Balac, Amedeo R. Vetrella, and Kay W. Axhausen. “Towards the Integration of Aerial Transportation in Urban Settings”. In: (2017). DOI: 10.3929/ethz-b-000175251.
- [21] Milos Balac, Amedeo R. Vetrella, and Kay W. Axhausen. “Towards the Integration of Aerial Transportation in Urban Settings”. In: *Transportation Research Board* (2018). DOI: 10.3929/ethz-b-000175251. URL: <http://hdl.handle.net/20.500.11850/175251> (visited on 05/23/2018).
- [22] Milos Balac et al. “Demand Estimation for Aerial Vehicles in Urban Settings”. In: *IEEE Intelligent Transportation Systems Magazine* 11.3 (23–2019), pp. 105–116. ISSN: 1939-1390, 1941-1197. DOI: 10.1109/MITS.2019.2919500. URL: <https://ieeexplore.ieee.org/document/8744222/> (visited on 04/15/2020).
- [23] Milos Balać, Amedeo R. Vetrella, and Kay W. Axhausen. “Towards the Integration of Aerial Transportation in Urban Settings”. In: (2018), 12 p. DOI: 10.3929/ETHZ-B-000193150. URL: <http://hdl.handle.net/20.500.11850/193150> (visited on 11/09/2020).
- [24] D Banister, S Watson, and C Wood. “Sustainable Cities: Transport, Energy, and Urban Form”. In: *Environment and Planning B: Planning and Design* 24.1 (Feb. 1, 1997), pp. 125–143. ISSN: 0265-8135. DOI: 10.1068/b240125. URL: <https://journals.sagepub.com/doi/abs/10.1068/b240125> (visited on 04/27/2019).
- [25] Hugo Barbosa et al. “Human Mobility: Models and Applications”. In: *Physics Reports* 734 (Mar. 2018), pp. 1–74. ISSN: 03701573. DOI: 10.1016/j.physrep.2018.01.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S037015731830022X> (visited on 06/10/2020).
- [26] Jaime Barceló and Jordi Casas. “Dynamic Network Simulation with AIMSUN”. In: *Simulation Approaches in Transportation Analysis*. Ed. by Ryuichi Kitamura and Maso Kuwahara. Vol. 31. Operations Research/Computer Science Interfaces Series. New York: Springer-Verlag, 2005, pp. 57–98. ISBN: 978-0-387-24108-1. DOI: 10.1007/0-387-24109-4\_3. URL: [http://link.springer.com/10.1007/0-387-24109-4\\_3](http://link.springer.com/10.1007/0-387-24109-4_3) (visited on 05/16/2020).
- [27] Jaume Barceló et al. “Parallelization of Microscopic Traffic Simulation for Att Systems Analysis”. In: *Equilibrium and Advanced Transportation Modelling*. Ed. by Patrice Marcotte and Sang Nguyen. Boston, MA: Springer US, 1998, pp. 1–26. ISBN: 978-1-4613-7638-5 978-1-4615-5757-9. DOI: 10.1007/978-1-4615-5757-9\_1. URL: [http://link.springer.com/10.1007/978-1-4615-5757-9\\_1](http://link.springer.com/10.1007/978-1-4615-5757-9_1) (visited on 06/08/2020).

- [28] M. Batty et al. “Smart Cities of the Future”. In: *The European Physical Journal Special Topics* 214.1 (Nov. 2012), pp. 481–518. ISSN: 1951-6355, 1951-6401. DOI: 10.1140/epjst/e2012-01703-3. URL: <http://link.springer.com/10.1140/epjst/e2012-01703-3> (visited on 03/28/2021).
- [29] Moshe Ben-Akiva. *Discrete Choice Analysis*. The MIT Press. URL: <https://mitpress.mit.edu/books/discrete-choice-analysis> (visited on 02/13/2019).
- [30] Moshe Ben-Akiva et al. “DynaMIT: A Simulation-Based System for Traffic Prediction”. In: *DACCORD Short Term Forecasting Workshop* (1998), p. 12.
- [31] Dimitri Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*. URL: [https://books.google.com/books/about/Constrained\\_Optimization\\_and\\_Lagrange\\_Mu.html?id=j6LiBQAAQBAJ](https://books.google.com/books/about/Constrained_Optimization_and_Lagrange_Mu.html?id=j6LiBQAAQBAJ) (visited on 04/28/2019).
- [32] Joschka Bischoff and Michal Maciejewski. “Simulation of City-Wide Replacement of Private Cars with Autonomous Taxis in Berlin”. In: *Procedia Computer Science* 83 (2016), pp. 237–244. ISSN: 18770509. DOI: 10.1016/j.procs.2016.04.121. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1877050916301442> (visited on 05/24/2018).
- [33] Marlon G. Boarnet. “A Broader Context for Land Use and Travel Behavior, and a Research Agenda”. In: *Journal of the American Planning Association* 77.3 (July 1, 2011), pp. 197–213. ISSN: 0194-4363. DOI: 10.1080/01944363.2011.593483. URL: <https://doi.org/10.1080/01944363.2011.593483> (visited on 12/10/2018).
- [34] Geoff Boeing. “OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks”. In: *Computers, Environment and Urban Systems* 65 (Sept. 2017), pp. 126–139. ISSN: 01989715. DOI: 10.1016/j.compenurbsys.2017.05.004. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0198971516303970> (visited on 05/13/2020).
- [35] Helena Bohman and Désirée Nilsson. “The Impact of Regional Commuter Trains on Property Values: Price Segments and Income”. In: *Journal of Transport Geography* 56 (Oct. 1, 2016), pp. 102–109. ISSN: 0966-6923. DOI: 10.1016/j.jtrangeo.2016.09.003. URL: <http://www.sciencedirect.com/science/article/pii/S0966692316300151> (visited on 03/05/2019).
- [36] Stephen P. Borgatti. “Centrality and Network Flow”. In: *Social Networks* 27.1 (2005), pp. 55–71. ISSN: 0378-8733. DOI: 10.1016/j.socnet.2004.11.008. URL: <http://www.sciencedirect.com/science/article/pii/S0378873304000693>.
- [37] Sanjay Bose. *An Introduction to Queuing Systems*. Jan. 2002. ISBN: 978-1-4613-4880-1. DOI: 10.1007/978-1-4615-0001-8.
- [38] Elmar Brockfeld, Peter Wagner, and Reinhart Kuehne. “Calibration and Validation of Microscopic Traffic Flow Models”. In: (2004), p. 15.

- [39] Arthur Brown and Wesley Harris. “A Vehicle Design and Optimization Model for On-Demand Aviation”. In: American Institute of Aeronautics and Astronautics, Jan. 8, 2018. ISBN: 978-1-62410-532-6. DOI: 10.2514/6.2018-0105. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-0105> (visited on 05/30/2018).
- [40] Jay R. Brown and Alfred L. Guiffrida. “Carbon Emissions Comparison of Last Mile Delivery versus Customer Pickup”. In: *International Journal of Logistics Research and Applications* 17.6 (Nov. 2, 2014), pp. 503–521. ISSN: 1367-5567, 1469-848X. DOI: 10.1080/13675567.2014.907397. URL: <http://www.tandfonline.com/doi/abs/10.1080/13675567.2014.907397> (visited on 07/02/2021).
- [41] D. J. Buckley. “A Semi-Poisson Model of Traffic Flow”. In: *Transportation Science* 2.2 (May 1968), pp. 107–133. ISSN: 0041-1655, 1526-5447. DOI: 10.1287/trsc.2.2.107. URL: <http://pubsonline.informs.org/doi/abs/10.1287/trsc.2.2.107> (visited on 02/14/2019).
- [42] Vishwanath Bulusu et al. “A Traffic Demand Analysis Method for Urban Air Mobility”. In: *IEEE Intelligent Transportation Systems Magazine* (2020).
- [43] Robert W Burchell and Naveed A Shad. “The Evolution of the Sprawl Debate in the United States”. In: (), p. 25.
- [44] Robert Cervero. “Going Beyond Travel-Time Savings”. In: *The World Bank* (2011), p. 47.
- [45] Robert Cervero and Kang-Li Wu. “Sub-Centring and Commuting: Evidence from the San Francisco Bay Area, 1980-90”. In: *Urban Studies* 35.7 (June 1, 1998), pp. 1059–1076. ISSN: 0042-0980. DOI: 10.1080/0042098984484. URL: <https://doi.org/10.1080/0042098984484> (visited on 04/27/2019).
- [46] Cy Chan et al. “Mobiliti: Scalable Transportation Simulation Using High-Performance Parallel Computing”. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. 2018 21st International Conference on Intelligent Transportation Systems (ITSC). Maui, HI: IEEE, Nov. 2018, pp. 634–641. ISBN: 978-1-72810-321-1 978-1-72810-323-5. DOI: 10.1109/ITSC.2018.8569397. URL: <https://ieeexplore.ieee.org/document/8569397/> (visited on 10/02/2019).
- [47] Victoria Chang, Pramod Chundury, and Marshini Chetty. “Spiders in the Sky: User Perceptions of Drones, Privacy, and Security”. In: ACM Press, 2017, pp. 6765–6776. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025632. URL: <http://dl.acm.org/citation.cfm?doid=3025453.3025632> (visited on 05/30/2018).
- [48] Daniel G. Chatman and Nicholas Klein. “Immigrants and Travel Demand in the United States: Implications for Transportation Policy and Future Research”. In: *Public Works Management & Policy* 13.4 (Apr. 2009), pp. 312–327. ISSN: 1087-724X, 1552-7549. DOI: 10.1177/1087724X09334633. URL: <http://journals.sagepub.com/doi/10.1177/1087724X09334633> (visited on 12/10/2018).



- [49] Charisma Choudhury et al. “Modeling Cooperative Lane-Changing and Forced Merging Behavior”. In: Jan. 1, 2007.
- [50] National Association of City Transportation Officials. *Multimodal Streets Serve More People*. Global Designing Cities Initiative. URL: <https://globaldesigningcities.org/publication/global-street-design-guide/defining-streets/multimodal-streets-serve-people/> (visited on 10/11/2018).
- [51] Lewis M. Clements and Kara M. Kockelman. “Economic Effects of Automated Vehicles”. In: *Transportation Research Record: Journal of the Transportation Research Board* 2606.1 (Jan. 2017), pp. 106–114. ISSN: 0361-1981, 2169-4052. DOI: 10.3141/2606-14. URL: <http://journals.sagepub.com/doi/10.3141/2606-14> (visited on 04/30/2021).
- [52] Regina R Clewlow et al. “Disruptive Transportation: The Adoption, Utilization, and Impacts of Ride-Hailing in the United States FOR MEDIA OR OTHER INQUIRIES”. In: (2017). URL: [file:///Users/pavan/Downloads/2017\\_UCD-ITS-RR-17-07.pdf](file:///Users/pavan/Downloads/2017_UCD-ITS-RR-17-07.pdf).
- [53] Reece A. Clothier et al. “Risk Perception and the Public Acceptance of Drones: Risk Perception and the Public Acceptance of Drones”. In: *Risk Analysis* 35.6 (June 2015), pp. 1167–1183. ISSN: 02724332. DOI: 10.1111/risa.12330. URL: <http://doi.wiley.com/10.1111/risa.12330> (visited on 05/30/2018).
- [54] Adam Cohen, Susan Shaheen, and E Farrar. “Urban Air Mobility: History, Ecosystem, Market Potential, and Challenges”. In: *IEEE Transactions on Intelligent Transportation Systems* (2021).
- [55] Marc M. Cohen. “The Vertiport as an Urban Design Problem”. In: World Aviation Congress & Exposition. Oct. 1, 1996. DOI: 10.4271/965523. URL: <https://www.sae.org/content/965523/> (visited on 05/06/2020).
- [56] Serdar Çolak, Antonio Lima, and Marta C. González. “Understanding Congested Travel in Urban Areas”. In: *Nature Communications* 7 (Mar. 2016), pp. 10793–10793. DOI: 10.1038/ncomms10793. URL: <http://www.nature.com/doi/10.1038/ncomms10793>.
- [57] Gabor Csardi, Tamas Nepusz, et al. “The Igraph Software Package for Complex Network Research”. In: *InterJournal, complex systems* 1695.5 (2006), pp. 1–9.
- [58] Saulo B. Cwerner. “Vertical Flight and Urban Mobilities: The Promise and Reality of Helicopter Travel”. In: *Mobilities* 1.2 (June 2006), pp. 191–215. DOI: 10.1080/17450100600726589. URL: <http://www.tandfonline.com/doi/abs/10.1080/17450100600726589>.
- [59] Carlos F. Daganzo. *Fundamentals of Transportation and Traffic Operations*. Emerald Group Publishing Limited, Sept. 1997. ISBN: 978-0-08-042785-0. DOI: 10.1108/9780585475301. URL: <http://www.emeraldinsight.com/doi/book/10.1108/9780585475301>.

- [60] Carlos F. Daganzo. “The Cell Transmission Model: A Dynamic Representation of Highway Traffic Consistent with the Hydrodynamic Theory”. In: *Transportation Research Part B: Methodological* 28.4 (1994), pp. 269–287.
- [61] Sanjoy Dasgupta et al. *Explainable K-Means and k-Medians Clustering*. Feb. 27, 2020. arXiv: 2002.12538 [cs, stat]. URL: <http://arxiv.org/abs/2002.12538> (visited on 07/20/2020).
- [62] Matthew Daskilewicz et al. “Progress in Vertiport Placement and Estimating Aircraft Range Requirements for eVTOL Daily Commuting”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-2884. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2884> (visited on 06/09/2020).
- [63] Hal Daume. *A Course in Machine Learning*. URL: <http://ciml.info/> (visited on 05/16/2019).
- [64] Ricardo A Daziano, Mauricio Sarrias, and Benjamin Leard. “Are Consumers Willing to Pay to Let Cars Drive for Them?: Analyzing Response to Autonomous Vehicles”. In: (), p. 34.
- [65] Adriaan de Graaff. “Aviation Safety, an Introduction”. In: *Air & Space Europe* 3.3 (May 1, 2001), pp. 203–205. ISSN: 1290-0958. DOI: 10.1016/S1290-0958(01)90095-4. URL: <http://www.sciencedirect.com/science/article/pii/S1290095801900954> (visited on 04/28/2019).
- [66] Mariusz Bojarski Davide Del and Testa Daniel Dworakowski Bernhard. “End to End Learning for Self-Driving Cars”. In: (), p. 9.
- [67] Daniel Delling et al. “Engineering Route Planning Algorithms”. In: *Algorithmics of Large and Complex Networks: Design, Analysis, and Simulation*. Ed. by Jürgen Lerner, Dorothea Wagner, and Katharina A. Zweig. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 117–139. ISBN: 978-3-642-02094-0. DOI: 10.1007/978-3-642-02094-0\_7. URL: [https://doi.org/10.1007/978-3-642-02094-0\\_7](https://doi.org/10.1007/978-3-642-02094-0_7) (visited on 06/02/2019).
- [68] Deloitte Insights. *The Elevated Future of Mobility*. 2019. URL: <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/energy-resources/di-the-elevated-future-of-mobility.pdf> (visited on 10/31/2020).
- [69] Coline Devin et al. *Deep Object-Centric Representations for Generalizable Robot Learning*. Aug. 14, 2017. arXiv: 1708.04225 [cs]. URL: <http://arxiv.org/abs/1708.04225> (visited on 09/07/2018).
- [70] D. N. Dewees. “The Effect of a Subway on Residential Property Values in Toronto”. In: *Journal of Urban Economics* 3.4 (Oct. 1, 1976), pp. 357–369. ISSN: 0094-1190. DOI: 10.1016/0094-1190(76)90035-8. URL: <http://www.sciencedirect.com/science/article/pii/0094119076900358> (visited on 03/05/2019).

- [71] David Dohan, Stefani Karp, and Brian Matejek. *K-Median Algorithms: Theory in Practice*. Jan. 2015.
- [72] Richard Dowling, Alexander Skabardonis, and Vassili Alexiadis. “Traffic Analysis Toolbox, Volume III: Guidelines for Applying Traffic Microsimulation Modeling Software”. In: *Transportation Research Board* (FHWA-HRT-04-040 June 1, 2004). Ed. by United States. Federal Highway Administration. Office of Operations. URL: <https://rosap.nhtl.bts.gov/view/dot/41797>.
- [73] Dr Graham Drozeski. “Vertical Take-off and Landing (VTOL): Emerging and Transformational Capabilities”. In: (), p. 8.
- [74] Jean Dubé, Marius Thériault, and François Des Rosiers. “Commuter Rail Accessibility and House Values: The Case of the Montreal South Shore, Canada, 1992–2009”. In: *Transportation Research Part A: Policy and Practice* 54 (Aug. 1, 2013), pp. 49–66. ISSN: 0965-8564. DOI: 10.1016/j.tra.2013.07.015. URL: <http://www.sciencedirect.com/science/article/pii/S0965856413001377> (visited on 03/05/2019).
- [75] Michael J. Duffy, Sean R. Wakayama, and Ryan Hupp. “A Study in Reducing the Cost of Vertical Flight with Electric Propulsion”. In: American Institute of Aeronautics and Astronautics, June 5, 2017. ISBN: 978-1-62410-508-1. DOI: 10.2514/6.2017-3442. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3442> (visited on 05/30/2018).
- [76] Gilles Duranton and Matthew A Turner. “The Fundamental Law of Road Congestion: Evidence from US Cities”. In: *American Economic Review* 101.6 (Oct. 2011), pp. 2616–2652. ISSN: 0002-8282. DOI: 10.1257/aer.101.6.2616. URL: <http://pubs.aeaweb.org/doi/10.1257/aer.101.6.2616> (visited on 10/11/2018).
- [77] Uber Elevate. “Fast-Forwarding to a Future of On-Demand Urban Air Transportation”. In: (2016). URL: <https://www.uber.com/elevate.pdf>.
- [78] Dimas Numan Fadhil, Rolf Moeckel, and Raoul Rothfeld. “GIS-Based Infrastructure Requirement Analysis for an Electric Vertical Take-off and Landing Vehicle-Based Transportation System”. In: *Transportation Research Procedia* 41 (2019), pp. 101–103. ISSN: 23521465. DOI: 10.1016/j.trpro.2019.09.020. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2352146519304314> (visited on 06/15/2021).
- [79] DN Fadhil. *A GIS-Based Analysis for Selecting Ground Infrastructure Locations for Urban Air Mobility*. 2018.
- [80] Daniel J. Fagnant and Kara M. Kockelman. “The Travel and Environmental Implications of Shared Autonomous Vehicles, Using Agent-Based Model Scenarios”. In: *Transportation Research Part C: Emerging Technologies* 40 (Mar. 2014), pp. 1–13. ISSN: 0968090X. DOI: 10.1016/j.trc.2013.12.001. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0968090X13002581> (visited on 05/24/2018).
- [81] Federal Aviation Administration. *AC 150/5390-2C - Heliport Design*. Apr. 2012.

- [82] L. T. Filotas. “Vortex Induced Helicopter Blade Loads and Noise”. In: *Journal of Sound and Vibration* 27.3 (Apr. 8, 1973), pp. 387–398. ISSN: 0022-460X. DOI: 10.1016/S0022-460X(73)80353-0. URL: <http://www.sciencedirect.com/science/article/pii/S0022460X73803530> (visited on 04/28/2019).
- [83] Chelsea Finn et al. *Deep Spatial Autoencoders for Visuomotor Learning*. Sept. 21, 2015. arXiv: 1509.06113 [cs]. URL: <http://arxiv.org/abs/1509.06113> (visited on 09/07/2018).
- [84] Fletcher Foti and Paul Waddell. “A Generalized Computational Framework for Accessibility: From the Pedestrian to the Metropolitan Scale”. In: (), p. 14.
- [85] Mengying Fu, Raoul Rothfeld, and Constantinos Antoniou. “Exploring Preferences for Transportation Modes in an Urban Air Mobility Environment: Munich Case Study”. In: *Transportation Research Record: Journal of the Transportation Research Board* 2673.10 (Oct. 2019), pp. 427–442. ISSN: 0361-1981, 2169-4052. DOI: 10.1177/0361198119843858. URL: <http://journals.sagepub.com/doi/10.1177/0361198119843858> (visited on 04/07/2020).
- [86] Juntao Gao et al. *Adaptive Traffic Signal Control: Deep Reinforcement Learning Algorithm with Experience Replay and Target Network*. May 8, 2017. arXiv: 1705.02755 [cs]. URL: <http://arxiv.org/abs/1705.02755> (visited on 10/15/2018).
- [87] I. Garcia-Dorado, D. G. Aliaga, and S. V. Ukkusuri. “Designing Large-Scale Interactive Traffic Animations for Urban Modeling: Designing Large-Scale Interactive Traffic Animations for Urban Modeling”. In: *Computer Graphics Forum* 33.2 (May 2014), pp. 411–420. ISSN: 01677055. DOI: 10.1111/cgf.12329. URL: <http://doi.wiley.com/10.1111/cgf.12329> (visited on 06/06/2018).
- [88] Laurie A. Garrow, Brian J. German, and Mohammad Ilbeigi. “Conceptual Models of Demand for Electric Propulsion Aircraft in Intra-Urban and Thin-Haul Markets”. In: Transportation Research Board 97th Annual Meeting Transportation Research Board. 2018. URL: <https://trid.trb.org/view/1495360> (visited on 06/18/2019).
- [89] Laurie A. Garrow, Mohammad Ilbeigi, and Ziran Chen. “Forecasting Demand for On Demand Mobility”. In: *17th AIAA Aviation Technology, Integration, and Operations Conference*. 0 vols. AIAA AVIATION Forum. American Institute of Aeronautics and Astronautics, June 5, 2017. DOI: 10.2514/6.2017-3280. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3280> (visited on 02/22/2019).
- [90] Laurie A. Garrow et al. “If You Fly It, Will Commuters Come? A Survey to Model Demand for eVTOL Urban Air Trips”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-2882. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2882> (visited on 11/18/2018).

- [91] Guido Gentile, Lorenzo Meschini, and Natale Papola. “Spillback Congestion in Dynamic Traffic Assignment: A Macroscopic Flow Model with Time-Varying Bottlenecks”. In: *Transportation Research Part B: Methodological* 41.10 (Dec. 2007), pp. 1114–1138. ISSN: 01912615. DOI: 10.1016/j.trb.2007.04.011. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0191261507000525> (visited on 05/20/2020).
- [92] Brian German et al. “Cargo Delivery in by Passenger eVTOL Aircraft: A Case Study in the San Francisco Bay Area”. In: *2018 AIAA Aerospace Sciences Meeting*. 2018 AIAA Aerospace Sciences Meeting. Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 8, 2018. ISBN: 978-1-62410-524-1. DOI: 10.2514/6.2018-2006. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2006> (visited on 06/09/2020).
- [93] N. Geroliminis and A. Skabardonis. “Identification and Analysis of Queue Spillovers in City Street Networks”. In: *IEEE Transactions on Intelligent Transportation Systems* 12.4 (Dec. 2011), pp. 1107–1115. DOI: 10.1109/TITS.2011.2141991.
- [94] Jonathan L. Gifford and William L. Garrison. “Airports and the Air Transportation System: Functional Refinements and Functional Discovery”. In: *Technological Forecasting and Social Change* 43.2 (Mar. 1, 1993), pp. 103–123. ISSN: 0040-1625. DOI: 10.1016/0040-1625(93)90011-U. URL: <http://www.sciencedirect.com/science/article/pii/004016259390011U> (visited on 02/14/2019).
- [95] Lillian Gipson. *NASA to Help Local Governments Plan for Advanced Air Mobility*. NASA. URL: <http://www.nasa.gov/aeroresearch/programs/iasp/aam/nasa-to-help-local-governments-plan-for-advanced-air-mobility> (visited on 07/02/2021).
- [96] Genevieve Giuliano and Susan Hanson. *The Geography of Urban Transportation*. 2017. 400 pp. ISBN: 978-1-4625-2969-8.
- [97] Genevieve Giuliano and Susan Hanson. *The Geography of Urban Transportation*. 2017. ISBN: 978-1-4625-2969-8.
- [98] P Gordon, H W Richardson, and H L Wong. “The Distribution of Population and Employment in a Polycentric City: The Case of Los Angeles”. In: *Environment and Planning A: Economy and Space* 18.2 (Feb. 1, 1986), pp. 161–173. ISSN: 0308-518X. DOI: 10.1068/a180161. URL: <https://doi.org/10.1068/a180161> (visited on 04/27/2019).
- [99] Dominik Grether and Kai Nagel. “Agent-Based Modelling of Air Transport Demand”. In: (), p. 22.
- [100] Bronwyn H Hall and Beethika Khan. “Adoption of New Technology”. In: (2002), p. 38.

- [101] Susan Handy. “Methodologies for Exploring the Link between Urban Form and Travel Behavior”. In: *Transportation Research Part D: Transport and Environment* 1.2 (Dec. 1996), pp. 151–165. ISSN: 13619209. DOI: 10.1016/S1361-9209(96)00010-7. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1361920996000107> (visited on 06/12/2018).
- [102] Susan Handy. “Smart Growth and the Transportation-Land Use Connection: What Does the Research Tell Us?” In: *International Regional Science Review* 28.2 (Apr. 1, 2005), pp. 146–167. ISSN: 0160-0176. DOI: 10.1177/0160017604273626. URL: <https://doi.org/10.1177/0160017604273626> (visited on 12/10/2018).
- [103] Susan L. Handy. “Understanding the Link Between Urban Form and Nonwork Travel Behavior”. In: *Journal of Planning Education and Research* 15.3 (Apr. 1, 1996), pp. 183–198. ISSN: 0739-456X. DOI: 10.1177/0739456X9601500303. URL: <https://doi.org/10.1177/0739456X9601500303> (visited on 04/27/2019).
- [104] Walter G. Hansen. “How Accessibility Shapes Land Use”. In: *Journal of the American Institute of Planners* 25.2 (May 1, 1959), pp. 73–76. ISSN: 0002-8991. DOI: 10.1080/01944365908978307. URL: <https://doi.org/10.1080/01944365908978307> (visited on 04/27/2019).
- [105] Kaiming He et al. *Deep Residual Learning for Image Recognition*. Dec. 10, 2015. arXiv: 1512.03385 [cs]. URL: <http://arxiv.org/abs/1512.03385> (visited on 10/20/2018).
- [106] Aude Hoeffleitner, Ryan Herring, and Alexandre M. Bayen. “Arterial Travel Time Forecast with Streaming Data: A Hybrid Approach of Flow Modeling and Machine Learning”. In: *Transportation Research Part B: Methodological* 46.9 (2012), pp. 1097–1122.
- [107] Bruce J. Holmes. “A Vision and Opportunity for Transformation of On-Demand Air Mobility”. In: *16th AIAA Aviation Technology, Integration, and Operations Conference*. 0 vols. AIAA AVIATION Forum. American Institute of Aeronautics and Astronautics, June 10, 2016. DOI: 10.2514/6.2016-3465. URL: <https://arc.aiaa.org/doi/10.2514/6.2016-3465> (visited on 02/22/2019).
- [108] Andreas Horni, Kai Nagel, and Kay Axhausen, eds. *Multi-Agent Transport Simulation MATSim*. London: Ubiquity Press, Aug. 2016. 618 pp. ISBN: 978-1-909188-75-4 978-1-909188-76-1 978-1-909188-77-8 978-1-909188-78-5. DOI: 10.5334/baw.
- [109] Yue Hu et al. *Impacts of Covid-19 Mode Shift on Road Traffic*. May 4, 2020. arXiv: 2005.01610 [physics]. URL: <http://arxiv.org/abs/2005.01610> (visited on 07/26/2020).
- [110] W. Huang et al. “Deep Architecture for Traffic Flow Prediction: Deep Belief Networks With Multitask Learning”. In: *IEEE Transactions on Intelligent Transportation Systems* 15.5 (Oct. 2014), pp. 2191–2201. ISSN: 1524-9050. DOI: 10.1109/TITS.2014.2311123.

- [111] Md. Shahadat Iqbal et al. “Development of Origin–Destination Matrices Using Mobile Phone Call Data”. In: *Transportation Research Part C: Emerging Technologies* 40 (Mar. 1, 2014), pp. 63–74. ISSN: 0968-090X. DOI: 10.1016/j.trc.2014.01.002. URL: <http://www.sciencedirect.com/science/article/pii/S0968090X14000059> (visited on 04/21/2019).
- [112] Martin Jaggi. “Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization”. In: *International Conference on Machine Learning*. International Conference on Machine Learning. Feb. 13, 2013, pp. 427–435. URL: <http://proceedings.mlr.press/v28/jaggi13.html> (visited on 02/14/2019).
- [113] Gareth James et al. *An Introduction to Statistical Learning: With Applications in R*. Springer Publishing Company, Incorporated, 2014. ISBN: 978-1-4614-7137-0.
- [114] Y. Jeong et al. “Supervised Weighting-Online Learning Algorithm for Short-Term Traffic Flow Prediction”. In: *IEEE Transactions on Intelligent Transportation Systems* 14.4 (Dec. 2013), pp. 1700–1707. ISSN: 1524-9050. DOI: 10.1109/TITS.2013.2267735.
- [115] Shan Jiang, Joseph Ferreira, and Marta C. González. “Clustering Daily Patterns of Human Activities in the City”. In: *Data Mining and Knowledge Discovery* 25.3 (Nov. 1, 2012), pp. 478–510. ISSN: 1573-756X. DOI: 10.1007/s10618-012-0264-z. URL: <https://doi.org/10.1007/s10618-012-0264-z> (visited on 04/21/2019).
- [116] W L Jin and H M Zhang. “Solving the Payne-Whitham Traffic Flow Model as a Hyperbolic System of Conservation Laws with Relaxation”. In: (), p. 24.
- [117] Akshat Kasliwal et al. “Role of Flying Cars in Sustainable Mobility”. In: *Nature Communications* 10.1 (Dec. 2019), p. 1555. ISSN: 2041-1723. DOI: 10.1038/s41467-019-09426-0. URL: <http://www.nature.com/articles/s41467-019-09426-0> (visited on 05/13/2019).
- [118] J.R. Kenworthy, P.W.G. Newman, and T.J. Lyons. “Urban Planning and Traffic Congestion”. In: *Urban Policy and Research* 7.2 (June 1989), pp. 67–80. ISSN: 0811-1146, 1476-7244. DOI: 10.1080/08111148908551389. URL: <http://www.tandfonline.com/doi/full/10.1080/08111148908551389> (visited on 06/11/2018).
- [119] Sanggyum Kim. “Design of the Adaptive Cruise Control Systems: An Optimal Control Approach”. In: (2012), p. 135.
- [120] Kittelson & Associates, Inc. et al. *Transit Capacity and Quality of Service Manual, Third Edition*. Washington, D.C.: Transportation Research Board, Aug. 26, 2013. ISBN: 978-0-309-28344-1. DOI: 10.17226/24766. URL: <https://www.nap.edu/catalog/24766> (visited on 10/16/2019).
- [121] Robert C Kloosterman and Sako Musterd. “The Polycentric Urban Region: Towards a Research Agenda”. In: (), p. 11.

- [122] Bradley Kloostera and Matthew J. Roorda. “Fully Autonomous Vehicles: Analyzing Transportation Network Performance and Operating Scenarios in the Greater Toronto Area, Canada”. In: *Transportation Planning and Technology* 42.2 (Feb. 17, 2019), pp. 99–112. ISSN: 0308-1060, 1029-0354. DOI: 10.1080/03081060.2019.1565159. URL: <https://www.tandfonline.com/doi/full/10.1080/03081060.2019.1565159> (visited on 07/02/2021).
- [123] Zafeiris Kokkinogenis et al. “Towards the Next-Generation Traffic Simulation Tools: A First Evaluation”. In: *Iberian Conference on Information Systems and Technologies* (2011), p. 14.
- [124] G Kotusevski and K A Hawick. “A Review of Traffic Simulation Software”. In: *Research Letters in the Information and Mathematical Sciences* (2009), p. 20.
- [125] Daniel Krajzewicz, Michael Bonert, and Peter Wagner. “The Open Source Traffic Simulation Package SUMO”. In: *RoboCup 2006*. RoboCup 2006. June 2006. URL: <https://elib.dlr.de/46740/>.
- [126] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Communications of the ACM* 60.6 (May 24, 2017), pp. 84–90. ISSN: 00010782. DOI: 10.1145/3065386. URL: <http://dl.acm.org/citation.cfm?doid=3098997.3065386> (visited on 10/20/2018).
- [127] Jonathan Lawson. “Flight Path: Taking to the Skies to Solve Congestion”. In: *Traffic Technology International* (June 2007). URL: <https://trid.trb.org/view.aspx?id=811269>.
- [128] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep Learning”. In: *Nature* 521.7553 (May 2015), pp. 436–444. ISSN: 1476-4687. DOI: 10.1038/nature14539. URL: <https://www.nature.com/articles/nature14539> (visited on 05/16/2019).
- [129] Robert V. Levine and Ara Norenzayan. “The Pace of Life in 31 Countries”. In: *Journal of Cross-Cultural Psychology* 30.2 (1999), pp. 178–205. DOI: 10.1177/0022022199030002003. eprint: <https://doi.org/10.1177/0022022199030002003>. URL: <https://doi.org/10.1177/0022022199030002003>.
- [130] Sheng Li, Maxim Egorov, and Mykel J. Kochenderfer. “Analysis of Fleet Management and Infrastructure Constraints in On-Demand Urban Air Mobility Operations”. In: *AIAA AVIATION 2020 FORUM*. AIAA AVIATION 2020 FORUM. VIRTUAL EVENT: American Institute of Aeronautics and Astronautics, June 15, 2020. ISBN: 978-1-62410-598-2. DOI: 10.2514/6.2020-2907. URL: <https://arc.aiaa.org/doi/10.2514/6.2020-2907> (visited on 07/26/2020).
- [131] Xiaoyuan Liang et al. *Deep Reinforcement Learning for Traffic Light Control in Vehicular Networks*. Mar. 29, 2018. arXiv: 1803.11115 [cs, stat]. URL: <http://arxiv.org/abs/1803.11115> (visited on 10/15/2018).



- [132] Eunha Lim and Hoyon Hwang. “The Selection of Vertiport Location for On-Demand Mobility and Its Application to Seoul Metro Area”. In: *International Journal of Aeronautical and Space Sciences* 20.1 (Mar. 2019), pp. 260–272. ISSN: 2093-274X, 2093-2480. DOI: 10.1007/s42405-018-0117-0. URL: <http://link.springer.com/10.1007/s42405-018-0117-0> (visited on 07/24/2020).
- [133] Allister Loder et al. “Understanding Traffic Capacity of Urban Networks”. In: *Scientific Reports* 9.1 (Dec. 2019), p. 16283. ISSN: 2045-2322. DOI: 10.1038/s41598-019-51539-5. URL: <http://www.nature.com/articles/s41598-019-51539-5> (visited on 04/16/2020).
- [134] Rémi Louf and Marc Barthelemy. “Modeling the Polycentric Transition of Cities”. In: (). DOI: 10.1103/PhysRevLett.111.198702. URL: <https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.111.198702>.
- [135] Y. Lv et al. “Traffic Flow Prediction With Big Data: A Deep Learning Approach”. In: *IEEE Transactions on Intelligent Transportation Systems* 16.2 (Apr. 2015), pp. 865–873. ISSN: 1524-9050. DOI: 10.1109/TITS.2014.2345663.
- [136] Sven Maerivoet and Bart De Moor. *Transportation Planning and Traffic Flow Models*. July 15, 2005. arXiv: physics/0507127. URL: <http://arxiv.org/abs/physics/0507127> (visited on 10/11/2018).
- [137] Sven Maerivoet and Bart De Moor. *Transportation Planning and Traffic Flow Models*. July 2005. arXiv: physics/0507127. URL: <http://arxiv.org/abs/physics/0507127> (visited on 10/11/2018).
- [138] C. Marchetti. “Anthropological Invariants in Travel Behavior”. In: *Technological Forecasting and Social Change* 47.1 (Sept. 1994), pp. 75–88. ISSN: 00401625. DOI: 10.1016/0040-1625(94)90041-8. URL: <https://linkinghub.elsevier.com/retrieve/pii/0040162594900418> (visited on 06/10/2020).
- [139] E Massaro. “The Car as an Ambient Sensing Platform”. In: *Proceedings of the IEEE* 105.3 (2017). DOI: 10.1109/JPROC.2016.2634938. URL: [http://senseable.mit.edu/papers/pdf/20170102\\_Massaroetal\\_CarAmbient\\_ProceedingsIEEE.pdf](http://senseable.mit.edu/papers/pdf/20170102_Massaroetal_CarAmbient_ProceedingsIEEE.pdf).
- [140] Samantha Masunaga. *A new generation of flying cars is taking to the air. But without the cars*. latimes.com. URL: <https://www.latimes.com/business/la-fi-flying-cars-20190222-story.html> (visited on 03/01/2019).
- [141] Klaus Mathis and Ariel David Steffen. “From Rational Choice to Behavioural Economics”. In: *European Perspectives on Behavioural Law and Economics*. Ed. by Klaus Mathis. Cham: Springer International Publishing, 2015, pp. 31–48. ISBN: 978-3-319-11634-1 978-3-319-11635-8. DOI: 10.1007/978-3-319-11635-8\_3. URL: [http://link.springer.com/10.1007/978-3-319-11635-8\\_3](http://link.springer.com/10.1007/978-3-319-11635-8_3) (visited on 02/14/2019).
- [142] Mckinsey & Company, Crown Consulting Inc., and NASA. “Urban Air Mobility Market Study”. Nov. 2018. URL: <https://www.nasa.gov/sites/default/files/atoms/files/uam-market-study-executive-summary-v2.pdf> (visited on 10/04/2019).

- [143] Metropolitan Transportation Commission and Association of Bay Area Governments. *Plan Bay Area 2040*. 2017. URL: [http://2040.planbayarea.org/files/2020-02/Travel\\_Modeling\\_PBA2040\\_Supplemental%20Report\\_7-2017.pdf](http://2040.planbayarea.org/files/2020-02/Travel_Modeling_PBA2040_Supplemental%20Report_7-2017.pdf) (visited on 10/30/2020).
- [144] Dimitris Milakis et al. “Is Micro-Mobility Sustainable? An Overview of Implications for Accessibility, Air Pollution, Safety, Physical Activity and Subjective Wellbeing”. In: Carey Curtis. *Handbook of Sustainable Transport*. Edward Elgar Publishing, 2020, pp. 180–189. ISBN: 978-1-78990-047-7. DOI: 10.4337/9781789900477.00030. URL: <https://www.elgaronline.com/view/edcoll/9781789900460/9781789900460.00030.xml> (visited on 07/02/2021).
- [145] Ronald T. Milam et al. “Closing the Induced Vehicle Travel Gap Between Research and Practice”. In: *Transportation Research Record: Journal of the Transportation Research Board* 2653 (Jan. 2017), pp. 10–16. ISSN: 0361-1981. DOI: 10.3141/2653-02. URL: <http://trrjournalonline.trb.org/doi/10.3141/2653-02> (visited on 10/11/2018).
- [146] Vicente Milanés and Steven E. Shladover. “Modeling Cooperative and Autonomous Adaptive Cruise Control Dynamic Responses Using Experimental Data”. In: *Transportation Research Part C: Emerging Technologies* 48 (Nov. 2014), pp. 285–300. ISSN: 0968090X. DOI: 10.1016/j.trc.2014.09.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0968090X14002447> (visited on 03/02/2021).
- [147] Mark D. Moore. “Misconceptions of Electric Aircraft and Their Emerging Aviation Markets”. In: *52nd Aerospace Sciences Meeting*. 0 vols. AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, Jan. 10, 2014. DOI: 10.2514/6.2014-0535. URL: <https://arc.aiaa.org/doi/10.2514/6.2014-0535> (visited on 02/22/2019).
- [148] Mark D. Moore and Kenneth H. Goodrich. “High Speed Mobility through On-Demand Aviation”. In: *2013 Aviation Technology, Integration, and Operations Conference*. 0 vols. AIAA AVIATION Forum. American Institute of Aeronautics and Astronautics, Aug. 8, 2013. DOI: 10.2514/6.2013-4373. URL: <https://arc.aiaa.org/doi/10.2514/6.2013-4373> (visited on 02/22/2019).
- [149] C. K. Moorthy and B. G. Ratcliffe. “Short Term Traffic Forecasting Using Time Series Methods”. In: *Transportation Planning and Technology* 12.1 (July 1, 1988), pp. 45–56. ISSN: 0308-1060. DOI: 10.1080/03081068808717359. URL: <https://doi.org/10.1080/03081068808717359> (visited on 05/15/2019).
- [150] Eric R. Mueller, Parmial H. Kopardekar, and Kenneth H. Goodrich. “Enabling Airspace Integration for High-Density On-Demand Mobility Operations”. In: *17th AIAA Aviation Technology, Integration, and Operations Conference*. 17th AIAA Aviation Technology, Integration, and Operations Conference. Denver, Colorado: American Institute of Aeronautics and Astronautics, June 5, 2017. ISBN: 978-1-62410-508-1. DOI:

- 10.2514/6.2017-3086. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3086> (visited on 02/22/2019).
- [151] Kai Nagel and Marcus Rickert. “Parallel Implementation of the TRANSIMS Micro-Simulation”. In: *Parallel Computing* 27.12 (Nov. 2001), pp. 1611–1639. ISSN: 01678191. DOI: 10.1016/S0167-8191(01)00106-5. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0167819101001065> (visited on 06/08/2020).
- [152] Paul R. Nail et al. “Threat Causes Liberals to Think like Conservatives”. In: *Journal of Experimental Social Psychology* 45.4 (July 2009), pp. 901–907. ISSN: 00221031. DOI: 10.1016/j.jesp.2009.04.013. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0022103109000948> (visited on 11/21/2018).
- [153] National Renewable Energy Laboratory. *Transportation Secure Data Center*. URL: <https://www.nrel.gov/tsdc> (visited on 11/02/2019).
- [154] Peter Newman. “Sustainable Transport for Sustainable Cities”. In: *Issues* 76 (Sept. 2006), p. 6. URL: <https://search.informit.com.au/documentSummary;dn=200611184;res=IELAPA> (visited on 05/07/2020).
- [155] Victoria C. Nneji et al. “Exploring Concepts of Operations for On-Demand Passenger Air Transportation”. In: American Institute of Aeronautics and Astronautics, June 5, 2017. ISBN: 978-1-62410-508-1. DOI: 10.2514/6.2017-3085. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3085> (visited on 04/11/2018).
- [156] Victoria C. Nneji et al. “Functional Requirements for Remotely Managing Fleets of On-Demand Passenger Aircraft”. In: American Institute of Aeronautics and Astronautics, Jan. 8, 2018. ISBN: 978-1-62410-524-1. DOI: 10.2514/6.2018-2007. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2007> (visited on 05/30/2018).
- [157] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. “Feature Visualization”. In: *Distill* 2.11 (Nov. 7, 2017). ISSN: 2476-0757. DOI: 10.23915/distill.00007. URL: <https://distill.pub/2017/feature-visualization> (visited on 10/26/2018).
- [158] Chris Olah et al. “The Building Blocks of Interpretability”. In: *Distill* 3.3 (Mar. 6, 2018). ISSN: 2476-0757. DOI: 10.23915/distill.00010. URL: <https://distill.pub/2018/building-blocks> (visited on 10/26/2018).
- [159] Stephen M. Omohundro and Stephen M. Omohundro. “Five Balltree Construction Algorithms”. In: (1989). URL: <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.91.8209>.
- [160] A. V. Oppenheim. *Applications of Digital Signal Processing*. 1978. URL: <http://adsabs.harvard.edu/abs/1978ph...bookR....0> (visited on 05/17/2019).
- [161] *Orlando Officials Hope the City Can Lead the Burgeoning 'air Taxi' Industry — Blogs*. URL: <https://www.orlandoweekly.com/Blogs/archives/2021/03/29/orlando-officials-pursing-multiple-air-taxi-projects-in-hopes-to-become-a-leader-in-the-industry> (visited on 07/02/2021).

- [162] Taner Osman. “Land Use Regulations and the Dispersion of the IT Industry in the San Francisco Bay Area”. In: *Papers in Regional Science* (Apr. 4, 2020). ISSN: 10568190. DOI: 10.1111/pirs.12532. URL: <http://doi.wiley.com/10.1111/pirs.12532> (visited on 06/01/2020).
- [163] R. K. Oswald, William T. Scherer, and Brian Lee Smith. “Traffic Flow Forecasting Using Approximate Nearest Neighbor Nonparametric Regression”. In: 2000.
- [164] Liviu Panait and Sean Luke. “Cooperative Multi-Agent Learning: The State of the Art”. In: *Autonomous Agents and Multi-Agent Systems* 11.3 (Nov. 2005), pp. 387–434. ISSN: 1387-2532, 1573-7454. DOI: 10.1007/s10458-005-2631-2. URL: <http://link.springer.com/10.1007/s10458-005-2631-2> (visited on 05/16/2019).
- [165] Byungkyu (Brian) Park and J. D. Schneeberger. “Microscopic Simulation Model Calibration and Validation: Case Study of VISSIM Simulation Model for a Coordinated Actuated Signal System”. In: *Transportation Research Record: Journal of the Transportation Research Board* 1856.1 (Jan. 2003), pp. 185–192. ISSN: 0361-1981, 2169-4052. DOI: 10.3141/1856-20. URL: <http://journals.sagepub.com/doi/10.3141/1856-20> (visited on 05/16/2020).
- [166] R M Pendyala, R Kitamura, and D V G Prasuna Reddy. “Application of an Activity-Based Travel-Demand Model Incorporating a Rule-Based Algorithm”. In: *Environment and Planning B: Planning and Design* 25.5 (Oct. 1, 1998), pp. 753–772. ISSN: 0265-8135. DOI: 10.1068/b250753. URL: <https://journals.sagepub.com/doi/abs/10.1068/b250753> (visited on 04/19/2019).
- [167] Pnina O. Plaut. “Non-Motorized Commuting in the US”. In: *Transportation Research Part D: Transport and Environment* 10.5 (Sept. 2005), pp. 347–356. ISSN: 13619209. DOI: 10.1016/j.trd.2005.04.002. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1361920905000179> (visited on 10/19/2018).
- [168] Nicholas G. Polson and Vadim O. Sokolov. “Deep Learning for Short-Term Traffic Flow Prediction”. In: *Transportation Research Part C: Emerging Technologies* 79 (June 1, 2017), pp. 1–17. ISSN: 0968-090X. DOI: 10.1016/j.trc.2017.02.024. URL: <http://www.sciencedirect.com/science/article/pii/S0968090X17300633> (visited on 05/13/2019).
- [169] Priyank Pradeep and Peng Wei. “Energy Efficient Arrival with RTA Constraint for Urban eVTOL Operations”. In: American Institute of Aeronautics and Astronautics, Jan. 8, 2018. ISBN: 978-1-62410-524-1. DOI: 10.2514/6.2018-2008. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2008> (visited on 05/30/2018).
- [170] *Press Release – FAA Statement–Federal vs. Local Drone Authority*. URL: [https://www.faa.gov/news/press\\_releases/news\\_story.cfm?newsId=22938](https://www.faa.gov/news/press_releases/news_story.cfm?newsId=22938) (visited on 07/02/2021).
- [171] *Principles of the Urban Sky*. World Economic Forum. URL: <https://www.weforum.org/reports/principles-of-the-urban-sky/> (visited on 07/02/2021).

- [172] Lisa M. PytlikZillig et al. “A Drone by Any Other Name: Purposes, End-User Trustworthiness, and Framing, but Not Terminology, Affect Public Support for Drones”. In: *IEEE Technology and Society Magazine* 37.1 (Mar. 2018), pp. 80–91. ISSN: 0278-0097. DOI: 10.1109/MTS.2018.2795121. URL: <http://ieeexplore.ieee.org/document/8307142/> (visited on 05/30/2018).
- [173] Sean Qian and Shuguan Yang. “What Do Autonomous Vehicles Mean to Traffic Congestion and Crash? Network Traffic Flow Modeling and Simulation for Autonomous Vehicles FINAL RESEARCH REPORT”. In: (). URL: [http://utc.ices.cmu.edu/utc/tier-one-reports/Qian1\\_TSETFinalReport.pdf](http://utc.ices.cmu.edu/utc/tier-one-reports/Qian1_TSETFinalReport.pdf).
- [174] Jayanthi Rajamani et al. “Assessing Impact of Urban Form Measures on Nonwork Trip Mode Choice After Controlling for Demographic and Level-of-Service Effects”. In: *Transportation Research Record* 1831.1 (Jan. 1, 2003), pp. 158–165. ISSN: 0361-1981. DOI: 10.3141/1831-18. URL: <https://doi.org/10.3141/1831-18> (visited on 04/27/2019).
- [175] Soora Rasouli and Harry Timmermans. “Activity-Based Models of Travel Demand: Promises, Progress and Prospects”. In: *International Journal of Urban Sciences* 18.1 (Jan. 2, 2014), pp. 31–60. ISSN: 1226-5934. DOI: 10.1080/12265934.2013.835118. URL: <https://doi.org/10.1080/12265934.2013.835118> (visited on 04/19/2019).
- [176] Colleene Reiche and Rohit Goyal. “Urban Air Mobility Market Study”. In: (2018). DOI: 10.7922/G2ZS2TRG. URL: <https://escholarship.org/uc/item/0fz0x1s2> (visited on 05/07/2021).
- [177] Joseph N. Robinson et al. “Development of a Methodology for Parametric Analysis of STOL Airpark Geo-Density”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-3054. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-3054> (visited on 07/24/2020).
- [178] Raoul Rothfeld et al. “Agent-Based Simulation of Urban Air Mobility”. In: *2018 Modeling and Simulation Technologies Conference*. 2018 Modeling and Simulation Technologies Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-551-7. DOI: 10.2514/6.2018-3891. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-3891> (visited on 11/28/2018).
- [179] Raoul Rothfeld et al. “Initial Analysis of Urban Air Mobility’s Transport Performance in Sioux Falls”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-2886. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-2886> (visited on 02/22/2019).

- [180] Raoul Rothfeld et al. “Potential Urban Air Mobility Travel Time Savings: An Exploratory Analysis of Munich, Paris, and San Francisco”. In: (June 2020).
- [181] Raoul Rothfeld et al. “Potential Urban Air Mobility Travel Time Savings: An Exploratory Analysis of Munich, Paris, and San Francisco”. In: *Sustainability* 13.4 (Feb. 2021), pp. 1–20. URL: <https://ideas.repec.org/a/gam/jsusta/v13y2021i4p2217-d501789.html>.
- [182] Mustapha Saidallah, Abdeslam El Fergougui, and Abdelbaki Elbelrhiti Elalaoui. “A Comparative Study of Urban Road Traffic Simulators”. In: *MATEC Web of Conferences* 81 (2016). Ed. by M. Figueira and Z. Guo, p. 05002. ISSN: 2261-236X. DOI: 10.1051/mateconf/20168105002. URL: <http://www.matec-conferences.org/10.1051/mateconf/20168105002> (visited on 09/07/2018).
- [183] Pratiksha Saxena, Dipti Singh, and Millie Pant, eds. *Problem Solving and Uncertainty Modeling through Optimization and Soft Computing Applications*: red. by Ivan Gianoccaro. Advances in Computational Intelligence and Robotics. IGI Global, 2016. ISBN: 978-1-4666-9885-7 978-1-4666-9886-4. DOI: 10.4018/978-1-4666-9885-7. URL: <http://services.igi-global.com/resolvedoi/resolve.aspx?doi=10.4018/978-1-4666-9885-7> (visited on 05/16/2020).
- [184] Jürgen Schmidhuber. “Deep Learning in Neural Networks: An Overview”. In: *Neural Networks* 61 (Jan. 1, 2015), pp. 85–117. ISSN: 0893-6080. DOI: 10.1016/j.neunet.2014.09.003. URL: <http://www.sciencedirect.com/science/article/pii/S0893608014002135> (visited on 05/16/2019).
- [185] R. Sengupta et al. “Smart Cities and Control [Technical Activities]”. In: *IEEE Control Systems Magazine* 35.6 (2015), pp. 20–21.
- [186] David A Senzig et al. “Sound Exposure Level Duration Adjustments in UAS Rotorcraft Noise Certification Tests”. In: (2018). DOI: 10.13140/RG.2.2.14083.76322. URL: <http://rgdoi.net/10.13140/RG.2.2.14083.76322> (visited on 07/02/2021).
- [187] Susan Shaheen and Adam Cohen. “The Potential Societal Barriers of Urban Air Mobility (UAM)”. In: (2018). DOI: 10.7922/G28C9TFR. URL: <https://escholarship.org/uc/item/7p69d2bg> (visited on 11/09/2020).
- [188] Michael Shamiyeh, Raoul Rothfeld, and Mirko Hornung. “A Performance Benchmark of Recent Personal Air Vehicle Concepts”. In: Sept. 9, 2018.
- [189] Dr Kapil Sheth. “Cascadia Urban Air Mobility Forum Feb. 28, 2020”. In: (2020), p. 8.
- [190] Julian Shun and Guy E Blelloch. “Ligra: A Lightweight Graph Processing Framework for Shared Memory”. In: *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 2013, pp. 135–146.
- [191] Christopher Silva et al. “VTOL Urban Air Mobility Concept Vehicles for Technology Development”. In: *2018 Aviation Technology, Integration, and Operations Conference*. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 2018. DOI: 10.2514/6.2018-3847.

- [192] ROBERT SINCLAIR. “Von Thünen and Urban Sprawl”. In: *Annals of the Association of American Geographers* 57.1 (Mar. 1, 1967), pp. 72–87. ISSN: 0004-5608. DOI: 10.1111/j.1467-8306.1967.tb00591.x. URL: <https://doi.org/10.1111/j.1467-8306.1967.tb00591.x> (visited on 04/26/2019).
- [193] Alexander Skabardonis, Pravin Varaiya, and Karl F. Petty. “Measuring Recurrent and Nonrecurrent Traffic Congestion”. In: *Transportation Research Record: Journal of the Transportation Research Board* 1856.1 (Jan. 2003), pp. 118–124. ISSN: 0361-1981, 2169-4052. DOI: 10.3141/1856-12. URL: <http://journals.sagepub.com/doi/10.3141/1856-12> (visited on 10/30/2020).
- [194] John Smart. *Our Amazing Aerial Future — How, When, and Why Air Taxis and Air Deliveries Will Change Our World*. Medium. Feb. 20, 2018. URL: <https://medium.com/@johnsmart/our-amazing-aerial-future-how-when-and-why-air-taxis-and-air-deliveries-will-change-our-world-2fc67d6b669> (visited on 05/23/2018).
- [195] Brian L Smith, Billy M Williams, and R Keith Oswald. “Comparison of Parametric and Nonparametric Models for Traffic Flow Forecasting”. In: *Transportation Research Part C: Emerging Technologies* 10.4 (Aug. 1, 2002), pp. 303–321. ISSN: 0968-090X. DOI: 10.1016/S0968-090X(02)00009-8. URL: <http://www.sciencedirect.com/science/article/pii/S0968090X02000098> (visited on 05/15/2019).
- [196] Min Geun Song and Gi Tae Yeo. “Analysis of the Air Transport Network Characteristics of Major Airports”. In: *The Asian Journal of Shipping and Logistics* 33.3 (2017), pp. 117–125. ISSN: 2092-5212. DOI: 10.1016/j.ajsl.2017.09.002. URL: <http://www.sciencedirect.com/science/article/pii/S209252121730041X>.
- [197] Raphael E. Stern et al. “Dissipation of Stop-and-Go Waves via Control of Autonomous Vehicles: Field Experiments”. In: *Transportation Research Part C: Emerging Technologies* 89 (Apr. 2018), pp. 205–221. ISSN: 0968090X. DOI: 10.1016/j.trc.2018.02.005. arXiv: 1705.01693. URL: <http://arxiv.org/abs/1705.01693> (visited on 02/14/2019).
- [198] Peter Stopher and John Stanley. *Introduction to Transport Policy*. Edward Elgar Publishing, 2014. URL: <https://EconPapers.repec.org/RePEc:elg:eebook:15102>.
- [199] Peter Stopher and John Stanley. *Introduction to Transport Policy: A Public Policy View*. Jan. 1, 2014. ISBN: 978 1 78195 246 7.
- [200] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. URL: <http://aima.cs.berkeley.edu/> (visited on 05/16/2019).
- [201] *Study: The Future of Vertical Mobility*. URL: <https://www.porsche-consulting.com/en/press/insights/detail/study-the-future-of-vertical-mobility/> (visited on 07/02/2021).

- [202] S. Sun, C. Zhang, and G. Yu. “A Bayesian Network Approach to Traffic Flow Forecasting”. In: *IEEE Transactions on Intelligent Transportation Systems* 7.1 (Mar. 2006), pp. 124–132. ISSN: 1524-9050. DOI: 10.1109/TITS.2006.869623. URL: <http://ieeexplore.ieee.org/document/1603558/> (visited on 05/15/2019).
- [203] Nida Syed et al. “Preliminary Considerations for ODM Air Traffic Management Based on Analysis of Commuter Passenger Demand and Travel Patterns for the Silicon Valley Region of California”. In: *17th AIAA Aviation Technology, Integration, and Operations Conference*. 17th AIAA Aviation Technology, Integration, and Operations Conference. Denver, Colorado: American Institute of Aeronautics and Astronautics, June 5, 2017. ISBN: 978-1-62410-508-1. DOI: 10.2514/6.2017-3082. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3082> (visited on 07/24/2020).
- [204] Technical Activities Division, Transportation Research Board, and National Academies of Sciences, Engineering, and Medicine. *Dynamic Traffic Assignment: A Primer*. Washington, D.C.: Transportation Research Board, July 6, 2011. ISBN: 978-0-309-43529-1. DOI: 10.17226/22872. URL: <https://www.nap.edu/catalog/22872> (visited on 02/14/2019).
- [205] *Texas Senate Bets on Urban Air Mobility as Transportation*. URL: <https://www.audacity.com/krlld/news/local/texas-senate-bets-on-urban-air-mobility-as-transportation> (visited on 07/02/2021).
- [206] David P. Thippavong et al. “Urban Air Mobility Airspace Integration Concepts and Considerations”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-3676. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-3676> (visited on 11/28/2018).
- [207] Trevor Thomas et al. “Not so Fast? Examining Neighborhood-Level Effects of Traffic Congestion on Job Access”. In: *Transportation Research Part A: Policy and Practice* 113 (July 1, 2018), pp. 529–541. ISSN: 0965-8564. DOI: 10.1016/j.tra.2018.04.015. URL: <http://www.sciencedirect.com/science/article/pii/S096585641730006X> (visited on 02/27/2019).
- [208] Sebastian Thrun et al. “Stanley: The Robot That Won the DARPA Grand Challenge”. In: *Journal of Field Robotics* 23.9 (Sept. 2006), pp. 661–692. ISSN: 15564959, 15564967. DOI: 10.1002/rob.20147. URL: <http://doi.wiley.com/10.1002/rob.20147> (visited on 02/14/2019).
- [209] Harry J. P. Timmermans and Junyi Zhang. “Modeling Household Activity Travel Behavior: Examples of State of the Art Modeling Approaches and Research Agenda”. In: *Transportation Research Part B: Methodological*. Modeling Household Activity Travel Behavior 43.2 (Feb. 1, 2009), pp. 187–190. ISSN: 0191-2615. DOI: 10.1016/j.trb.2008.06.004. URL: <http://www.sciencedirect.com/science/article/pii/S019126150800074X> (visited on 04/19/2019).



- [210] Tomer Toledo et al. “Microscopic Traffic Simulation: Models and Application”. In: *Simulation Approaches in Transportation Analysis*. New York: Springer-Verlag, 2005, pp. 99–130. DOI: 10.1007/0-387-24109-4\_4. URL: [http://link.springer.com/10.1007/0-387-24109-4\\_4](http://link.springer.com/10.1007/0-387-24109-4_4).
- [211] Kenneth Train. “Discrete Choice Methods with Simulation”. In: (), p. 388.
- [212] Martin Treiber and Arne Kesting. “An Open-Source Microscopic Traffic Simulator”. In: *IEEE INTELLIGENT TRANSPORTATION SYSTEMS MAGAZINE* 2.3 (2010), pp. 6–13. URL: <https://arxiv.org/pdf/1012.4913.pdf>.
- [213] Martin Treiber and Arne Kesting. *Traffic Flow Dynamics: Data, Models and Simulation*. Springer Berlin Heidelberg, 2013. ISBN: 978-3-642-32459-8.
- [214] Yu-Hsin Tsai. “Quantifying Urban Form: Compactness versus ‘Sprawl’”. In: *Urban Studies* 42.1 (Jan. 1, 2005), pp. 141–161. ISSN: 0042-0980. DOI: 10.1080/0042098042000309748. URL: <https://doi.org/10.1080/0042098042000309748> (visited on 03/02/2019).
- [215] Uber, Inc. *Uber Air Vehicle Mission and Requirements*. URL: <https://s3.amazonaws.com/uber-static/elevate/Summary+Mission+and+Requirements.pdf> (visited on 06/18/2019).
- [216] NASA Urban Air Mobility. *Uam-Market-Study-Executive-Summary-Pr.Pdf*. URL: <https://www.nasa.gov/sites/default/files/atoms/files/uam-market-study-executive-summary-pr.pdf> (visited on 04/28/2019).
- [217] OAR US EPA. *Fast Facts on Transportation Greenhouse Gas Emissions*. US EPA. Aug. 25, 2015. URL: <https://www.epa.gov/greenvehicles/fast-facts-transportation-greenhouse-gas-emissions> (visited on 04/28/2019).
- [218] James E. Vance. *Capturing the Horizon: The Historical Geography of Transportation since the Transportation Revolution of the Sixteenth Century*. New York: Harper & Row, 1986. xv, 656 p. ISBN: 978-0-06-046805-7. URL: <https://catalog.hathitrust.org/Record/000589768> (visited on 02/14/2019).
- [219] Parker D Vascik and R John Hansman. “Development of Vertiport Capacity Envelopes and Analysis of Their Sensitivity to Topological and Operational Factors”. In: *AIAA Scitech 2019 Forum*. 2019, p. 0526.
- [220] Parker D. Vascik and R John Hansman. “Evaluation of Key Operational Constraints Affecting On-Demand Mobility for Aviation in the Los Angeles Basin: Ground Infrastructure, Air Traffic Control and Noise”. In: American Institute of Aeronautics and Astronautics, June 5, 2017. ISBN: 978-1-62410-508-1. DOI: 10.2514/6.2017-3084. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3084> (visited on 05/30/2018).

- [221] Parker D. Vascik and R. John Hansman. “Development of Vertiport Capacity Envelopes and Analysis of Their Sensitivity to Topological and Operational Factors”. In: *AIAA Scitech 2019 Forum*. AIAA Scitech 2019 Forum. San Diego, California: American Institute of Aeronautics and Astronautics, Jan. 7, 2019. ISBN: 978-1-62410-578-4. DOI: 10.2514/6.2019-0526. URL: <https://arc.aiaa.org/doi/10.2514/6.2019-0526> (visited on 11/09/2020).
- [222] Parker D. Vascik and R. John Hansman. “Scaling Constraints for Urban Air Mobility Operations: Air Traffic Control, Ground Infrastructure, and Noise”. In: *2018 Aviation Technology, Integration, and Operations Conference*. American Institute of Aeronautics and Astronautics. DOI: 10.2514/6.2018-3849. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2018-3849> (visited on 04/28/2019).
- [223] Durgesh Vikram, Partha Chakroborty, and Sanjay Mittal. “Exploring the Behavior of LWR Continuum Models of Traffic Flow in Presence of Shock Waves”. In: *Procedia - Social and Behavioral Sciences*. 2nd Conference of Transportation Research Group of India (2nd CTRG) 104 (Dec. 2, 2013), pp. 412–421. ISSN: 1877-0428. DOI: 10.1016/j.sbspro.2013.11.134. URL: <http://www.sciencedirect.com/science/article/pii/S1877042813045254> (visited on 02/14/2019).
- [224] *Voom’s Helicopter Commuting Service Launches in Mexico City*. URL: <https://www.airbus.com/newsroom/news/en/2018/03/voom-s-helicopter-commuting-service-launches-in-mexico-city.html> (visited on 02/14/2019).
- [225] Paul Waddell. “Integrated Land Use and Transportation Planning and Modelling: Addressing Challenges in Research and Practice”. In: *Transport Reviews* 31.2 (Mar. 2011), pp. 209–229. ISSN: 0144-1647, 1464-5327. DOI: 10.1080/01441647.2010.525671. URL: <http://www.tandfonline.com/doi/abs/10.1080/01441647.2010.525671> (visited on 10/27/2020).
- [226] Paul Waddell. “UrbanSim Modeling Urban Development for Land Use, Transportation, and Environmental Planning”. In: *Journal of the American Planning Association* (2002). URL: <http://www.tandfonline.com/doi/pdf/10.1080/01944360208976274?needAccess=true>.
- [227] Paul Waddell et al. *An Integrated Pipeline Architecture for Modeling Urban Land Use, Travel Demand, and Traffic Assignment*. Feb. 7, 2018. arXiv: 1802.09335 [cs]. URL: <http://arxiv.org/abs/1802.09335> (visited on 02/14/2019).
- [228] Paul Waddell et al. *Architecture for Modular Microsimulation of Real Estate Markets and Transportation*. June 29, 2018. arXiv: 1807.01148 [cs]. URL: <http://arxiv.org/abs/1807.01148> (visited on 10/16/2020).
- [229] Rashid A. Waraich et al. “Performance Improvements for Large-Scale Traffic Simulation in MATSim”. In: *Computational Approaches for Urban Environments*. Ed. by Marco Helbich, Jamal Jokar Arsanjani, and Michael Leitner. Cham: Springer International Publishing, 2015, pp. 211–233. ISBN: 978-3-319-11468-2 978-3-319-11469-9. DOI:

- 10.1007/978-3-319-11469-9\_9. URL: [http://link.springer.com/10.1007/978-3-319-11469-9\\_9](http://link.springer.com/10.1007/978-3-319-11469-9_9) (visited on 05/14/2020).
- [230] Tom Wenzel et al. “Travel and Energy Implications of Ridesourcing Service in Austin, Texas”. In: *Transportation Research Part D: Transport and Environment* 70 (May 2019), pp. 18–34. ISSN: 13619209. DOI: 10.1016/j.trd.2019.03.005. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1361920918309878> (visited on 07/02/2021).
- [231] Cathy Wu et al. “Cellpath: Fusion of Cellular and Traffic Sensor Data for Route Flow Estimation via Convex Optimization”. In: *Transportation Research Procedia*. 21st International Symposium on Transportation and Traffic Theory Kobe, Japan, 5-7 August, 2015 7 (Jan. 1, 2015), pp. 212–232. ISSN: 2352-1465. DOI: 10.1016/j.trpro.2015.06.012. URL: <http://www.sciencedirect.com/science/article/pii/S2352146515000800> (visited on 05/09/2019).
- [232] Cathy Wu et al. *Flow: Architecture and Benchmarking for Reinforcement Learning in Traffic Control*. Oct. 15, 2017. arXiv: 1710.05465 [cs]. URL: <http://arxiv.org/abs/1710.05465> (visited on 08/15/2018).
- [233] Zhiqiang Wu and Yu Zhang. “Network Design, Performance Analysis, and Potential Demand Exploration of eVTOL On-Demand Service for Urban Air Mobility”. In: (2020).
- [234] Huazhe Xu et al. “End-to-End Learning of Driving Models from Large-Scale Video Datasets”. In: IEEE, July 2017, pp. 3530–3538. ISBN: 978-1-5386-0457-1. DOI: 10.1109/CVPR.2017.376. URL: <http://ieeexplore.ieee.org/document/8099859/> (visited on 04/26/2018).
- [235] Min Xue et al. “Fe<sup>3</sup>: An Evaluation Tool for Low-Altitude Air Traffic Operations”. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-3848. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-3848> (visited on 09/20/2019).
- [236] QI Yang and Haris N. Koutsopoulos. “A Microscopic Traffic Simulator for Evaluation of Dynamic Traffic Management Systems”. In: *Transportation Research Part C: Emerging Technologies* 4.3 (June 1996), pp. 113–129. DOI: 10.1016/S0968-090X(96)00006-X. URL: <https://www.sciencedirect.com/science/article/pii/S0968090X9600006X>.
- [237] Pavan Yedavalli. “Designing and Simulating Urban Air Mobility Vertiport Networks under Land Use Constraints”. In: *Transportation Research Board* (2021).
- [238] Pavan Yedavalli, Krishna Kumar, and Paul Waddell. “Microsimulation Analysis for Network Traffic Assignment (MANTA) at Metropolitan-Scale for Agile Transportation Planning”. In: (2020).

- [239] Pavan Yedavalli, Krishna Kumar, and Paul Waddell. *Microsimulation Analysis for Network Traffic Assignment (MANTA) at Metropolitan-Scale for Agile Transportation Planning*. July 4, 2020. arXiv: 2007.03614 [physics]. URL: <http://arxiv.org/abs/2007.03614> (visited on 07/20/2020).
- [240] Pavan Yedavalli, Krishna Kumar, and Paul Waddell. “Microsimulation Analysis for Network Traffic Assignment (MANTA) at Metropolitan-Scale for Agile Transportation Planning”. In: *Transportmetrica A: Transport Science* (June 11, 2021), pp. 1–22. ISSN: 2324-9935, 2324-9943. DOI: 10.1080/23249935.2021.1936281. URL: <https://www.tandfonline.com/doi/full/10.1080/23249935.2021.1936281> (visited on 06/14/2021).
- [241] Pavan Yedavalli and Jessie Mooberry. *An Assessment of Public Perception of Urban Air Mobility*. Feb. 2019. URL: <https://www.airbusutm.com/uam-resources-community-perception>.
- [242] Pavan Yedavalli and Jessie Mooberry. “An Assessment of Public Perception of Urban Air Mobility”. In: ().
- [243] Pavan Yedavalli et al. “Assessing the Value of Urban Air Mobility through Metropolitan-Scale Microsimulation: A Case Study of the San Francisco Bay Area”. In: *AIAA Aviation 2021* (2021).
- [244] Shannon Zelinski. “Operational Analysis of Vertiport Surface Topology”. In: *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)* (2020), pp. 1–10. DOI: 10.1109/DASC50938.2020.9256794.
- [245] Lun Zhang et al. “An Improved K-Nearest Neighbor Model for Short-Term Traffic Flow Prediction”. In: *Procedia - Social and Behavioral Sciences*. Intelligent and Integrated Sustainable Multimodal Transportation Systems Proceedings from the 13th COTA International Conference of Transportation Professionals (CICTP2013) 96 (Nov. 6, 2013), pp. 653–662. ISSN: 1877-0428. DOI: 10.1016/j.sbspro.2013.08.076. URL: <http://www.sciencedirect.com/science/article/pii/S1877042813022027> (visited on 05/15/2019).
- [246] Richard Zhang, Phillip Isola, and Alexei A. Efros. *Colorful Image Colorization*. Mar. 28, 2016. arXiv: 1603.08511 [cs]. URL: <http://arxiv.org/abs/1603.08511> (visited on 11/06/2018).
- [247] B. Zhao et al. “Agent-Based Model (ABM) for City-Scale Traffic Simulation: A Case Study on San Francisco”. In: *International Conference on Smart Infrastructure and Construction 2019 (ICSIC)*. International Conference on Smart Infrastructure and Construction 2019 (ICSIC). Cambridge, UK: ICE Publishing, Jan. 2019, pp. 203–212. ISBN: 978-0-7277-6466-9 978-0-7277-6467-6. DOI: 10.1680/icsic.64669.203. URL: <https://www.icevirtuallibrary.com/doi/10.1680/icsic.64669.203> (visited on 02/17/2021).